arXiv:1403.6212v1 [stat.ME] 25 Mar 2014

Selectable Factor Extraction in High Dimensions

Yiyuan She
Department of Statistics, Florida State University

This paper studies simultaneous rank reduction and variable selection in the vector regression setting. A novel selectable reduced rank regression is proposed for interpretable factor construction from a guaranteed small set of input features. The formulation provides a multivariate framework to pursue multiple sparse loading vectors jointly. Finite-sample studies reveal its benefit in predictive learning. Fast but simple-to-implement algorithms are developed and allow for essentially any sparsity-promoting penalties of practical interest. The technique can be adapted for rank constrained variable screening in ultrahigh dimensional applications. A nonasymptotic predictive information criterion is proposed for data-dependent parameter tuning, which can achieve the optimal prediction error rate. Experiments demonstrate that feature extraction and selection can be jointly achieved by simultaneous variable selection and projection.

1 Introduction

Modern statistical applications often involve a large number of variables, and the principal component analysis (PCA), dating back to Hotelling (1933) is a popular means of dimension reduction in real world data analysis. It offers an optimal low-dimensional projection subspace to view and analyze the data points, and is computationally efficient for large data. In the presence of response variables, PCA can be extended to the reduced rank regression (RRR) (Anderson, 1951; Izenman, 1975) for supervised learning. The rank reduction based regularization is very effective in multivariate data analysis (Reinsel and Velu, 1998).

Unfortunately, the construction of principal components in PCA (or factors in RRR) typically involves *all* dimensions, whereas in real-life high dimensional

data there may exist lots of nuisance variables. Because none of these irrelevant variables can be discarded and the important ones are hard to discern, sparse loading vectors are much more preferred practically. Moreover, recent theoretical studies (e.g., Johnstone and Lu (2009), Paul (2007)) have revealed that the ordinary PCA easily fails in moderate or high dimensions, in terms of loading vector recovery. A large class of sparse PCA have been put forward on the basis of sparsity-enforcing penalization, including, for example, Jolliffe et al. (2003), Zou et al. (2006), Shen and Huang (2008), Witten et al. (2009), Johnstone and Lu (2009), Journée et al. (2010), Ma (2013), among many others. On the other hand, most of these sparse PCA algorithms do not guarantee *joint* sparsity, optimality or orthogonality in seeking multiple loading vectors. Our goal is to construct a few best explanatory factors from a (guaranteed) reduced set of variables in a multivariate manner. This can also be used for factor-driven variable screening to tackle the challenge of ultrahigh dimensionality in genetics and computer vision applications.

This paper proposes and studies **selectable** factor extraction in the supervised RRR and the unsupervised PCA setups. Bunea et al. (2012) and Chen and Huang (2012) are perhaps the most relevant works in the literature. However, many theoretical and computational questions still remain open for simultaneous variable selection and rank reduction. For example, even concentrating on the error rate (see Section 2), there are no finite-sample results available in Chen and Huang (2012), and we show that the parameter choice and the rate established in Bunea et al. (2012) are far from being optimal. Computationally, a fast but simple-to-implement algorithm is in great need to meet the big data challenge; in particular, in consideration of some recent advances (e.g., Zhang and Zhang (2012)), nonconvex penalties are better to be allowed. Finally, the parameter tuning in such a joint regularization is non-trivial and should be addressed in finite samples.

The rest of the paper is organized as follows. Section 2 studies sparsity and selectability and formulates a framework for selectable factor extraction. We perform nonasymptotic analysis to reveal tight prediction error rate and propose a computational algorithm to handle nonconvex penalties. A sparse variant applies to the rank-r sparse PCA. Section 3 studies the model selection problem, where we prove a finite-sample oracle inequality under no constraints on the design or signal, and then propose a predictive information criterion (PIC). In Section 4, we develop rank constrained variable screening to facilitate the computation in ultrahigh dimensions and advocate a combined selectable/sparse reduced rank strategy in applications. Section 5 presents some data examples. All technical details are left to Section 6.

2 Simultaneous rank reduction and variable selection

2.1 Sparsity and selectability

We begin with the sparse PCA problem to illustrate the motivation of selectability. Suppose the data contain n observations in p-dimensional space, represented by a matrix $\boldsymbol{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_p] \in \mathbb{R}^{n\times p}$. Let $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T = \sum d_k \boldsymbol{u}_k \boldsymbol{v}_k^T$ be its SVD with all singular values d_k arranged in descending order. In PCA, the top r *PC directions* (or *loading vectors*) are then formed by the first r right singular vectors $[\boldsymbol{v}_1, \cdots, \boldsymbol{v}_r] =: \boldsymbol{V}_r$. PCA can be formulated in various ways; in particular, it can be viewed as a special case of the RRR which minimizes $\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2$ subject to $\text{rank}(\boldsymbol{B}) \le r$, where $\boldsymbol{Y}$ is the given response matrix.

The plain PCA has however certain drawbacks; one conventional criticism is that the PC construction $\boldsymbol{X}\boldsymbol{v}_k$ typically involves all input variables, and loses interpretability in high-dimensional data applications. Sparse PCA as a significant improvement has thus attracted much research attention recently. In seek of sparse loading vectors, most sparse PCA algorithms proceed in a *sequential* fashion. For example, an ℓ_0-type rank-1 approximation problem can be defined by $\min \|\boldsymbol{X} - d\boldsymbol{u}\boldsymbol{v}^T\|_F^2 + \lambda\|\boldsymbol{v}\|_0$ s.t. $\boldsymbol{u} \in \mathbb{R}^n : \|\boldsymbol{u}\|_2^2 = 1, \boldsymbol{v} \in \mathbb{R}^p : \|\boldsymbol{v}\|_2^2 = 1, d \ge 0$, or equivalently, $\max_{\boldsymbol{v}\in\mathbb{R}^p} \boldsymbol{v}^T(\boldsymbol{X}^T\boldsymbol{X})\boldsymbol{v} - \lambda\|\boldsymbol{v}\|_0$ s.t. $\|\boldsymbol{v}\|_2^2 = 1$, after evaluating the optimal d (as a function of $\boldsymbol{v}$ and $\boldsymbol{u}$). The remaining sparse PCA solutions $\boldsymbol{v}_k$ $(2 \le k \le r)$ can be obtained by repeatedly solving the rank-1 optimization after substituting the **deflated** matrix $\boldsymbol{X}_k = \boldsymbol{X}_{k-1} - d_{k-1}\boldsymbol{u}_{k-1}\boldsymbol{v}_{k-1}^T$ for $\boldsymbol{X}$. The ℓ_0-norm penalty can be replaced by, say, the ℓ_1-penalty, the SCAD penalty, or the fused ℓ_1 penalty in various applications. Even with the ℓ_1-relaxation, the problem is nonconvex.

Unfortunately, sparse PCA may not meet all application needs. (i) First, although each $\boldsymbol{v}_k$ is sparse, in constructing the top r principal directions, a number of input variables may be employed in total. The sequential rank-1 optimization scheme is incapable of enforcing joint sparsity. (ii) Secondly, there is no guarantee that the sequentially obtained $\boldsymbol{V}_r = [\boldsymbol{v}_1, \cdots \boldsymbol{v}_r]$, based on the rank-1 approximation and matrix deflation, satisfies certain desired properties such as joint orthogonality, joint optimality, and joint sparsity. It is worth pointing out that the deflation step has actually many possible choices and needs justification as well (Mackey, 2008). (iii) Finally, the optimization of each $\boldsymbol{v}_k$ $(1 \le k \le r)$ is conducted in p-dimensional space, resulting in a high computational burden especially when p

is very large. This may however be unnecessary due to the sparsity in $\boldsymbol{v}_k$. Therefore, in this paper, we will study multiple factor construction from a synthesis and multivariate perspective.

To state a general framework for joint variable selection and projection, we assume a response matrix $\boldsymbol{Y} \in \mathbb{R}^{n\times m}$ is available, in addition to the predictor matrix $\boldsymbol{X} \in \mathbb{R}^{n\times p}$, and write the coefficient matrix as $\boldsymbol{B} = [\boldsymbol{b}_1, \cdots, \boldsymbol{b}_p]^T = [b_{j,k}] \in \mathbb{R}^{p\times m}$. From the previous discussion, two types of regularization are desired for dimension reduction in high-dimensional applications, and can be achieved by minimizing the following objection function

$$\frac{1}{2}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2 + \frac{\lambda_1^2}{2}\|\boldsymbol{B}\|_{2,0} + \frac{\lambda_2^2}{2}\text{rank}(\boldsymbol{B}), \tag{1}$$

where $\|\boldsymbol{B}\|_{2,0} := \sum_j 1_{\boldsymbol{b}_j \neq \boldsymbol{0}}$, and λ_1 and λ_2 are two regularization parameters. The rank penalty is to bring low rankness for new feature extraction, while the group ℓ_0 penalty is to induce row-sparsity on $\boldsymbol{B}$ for feature selection, in which way feature extraction and selection can be simultaneously achieved. In this paper we favor the *rank constrained* variable selection defined by

$$\textbf{SEL-RRR:}\quad \min \frac{1}{2}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2 + \sum P(\|\boldsymbol{b}_j\|_2; \lambda) \quad s.t. \quad \text{rank}(\boldsymbol{B}) \leq r, \tag{2}$$

where P is a sparsity-promoting penalty possibly nonconvex. (When $\boldsymbol{X}$ and $\boldsymbol{Y}$ are not centered (columnwise), an intercept matrix should be included and is usually not subject to any constraint or penalty.) We call (2) the *selectable reduced-rank regression* (SEL-RRR). As will be shown in the next subsection, (2) facilitates the algorithm design in comparison to (1). A doubly constrained form of SEL-RRR will be studied in Section 4. Another example is the rank constrained group lasso (RCGL) (Bunea et al., 2011) which uses the ℓ_1 penalty for P. In addition, we propose the selectable principal component analysis (**SEL-PCA**) as a particular instance of (2)

$$\textbf{SEL-PCA:}\quad \min \frac{1}{2}\|\boldsymbol{X}^T - \boldsymbol{B}\|_F^2 + \sum P(\|\boldsymbol{b}_j\|_2; \lambda) \quad s.t. \quad \text{rank}(\boldsymbol{B}) \leq r. \tag{3}$$

The group regularization in SEL-PCA is different from that in sparse PCA. The penalty in group form can get rid of nuisance dimensions to give a reduced rank approximation with fewest possible input features. The special case of $r = 1$ corresponds to the rank-1 sparse PCA problem discussed previously. Another **SPA-PCA** variant is proposed in Section 2.2 to give a multivariate characterization of

sparse PCA. Our formulation provides a multivariate framework for simultaneously seeking top r loading vectors with elementwise or joint sparsity.

There are many choices for sparsity-promoting penalties, e.g., the group ℓ_1 (Yuan and Lin, 2006): $\sum_{j=1}^{p} \lambda_j \|\boldsymbol{b}_j\|_2$, or $\lambda \|\boldsymbol{B}\|_{2,1}$ when $\lambda_j = \lambda$, and the group ℓ_0: $\sum_{j=1}^{p} (\lambda_j^2/2) 1_{\boldsymbol{b}_j \neq \boldsymbol{0}}$, or $(\lambda^2/2)\|\boldsymbol{B}\|_{2,0}$ when $\lambda_j = \lambda$, as mentioned previously. Group SCAD, group l_{p}, and group $\ell_0 + \ell_2$ offer some fusion penalties. Computational algorithms will be developed to allow for all such practically used penalties (not necessarily convex).

After obtaining the estimate $\hat{\boldsymbol{B}}$, we can project the observed data points onto a low-dimensional subspace and construct new orthogonal features/factors. A direct way is to extract the range or column space of $\hat{\boldsymbol{B}}$. For example, one can use the reduced form SVD of $\hat{\boldsymbol{B}}$: $\hat{\boldsymbol{B}} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T$, to make a new model matrix $\boldsymbol{Z} = \boldsymbol{X}\boldsymbol{U}$ (or $\boldsymbol{Z} = \boldsymbol{X}\boldsymbol{U}\boldsymbol{D} = \boldsymbol{X}\hat{\boldsymbol{B}}\boldsymbol{V}$), which provides r factors with sparse loadings. In implementation, $\boldsymbol{Z}$ can be obtained more efficiently from the QR decomposition. We refer to it as the *Type-I extraction*. It can be used for parameter tuning. Alternatively, one can decompose $\boldsymbol{X}\hat{\boldsymbol{B}}$: Let $\hat{\boldsymbol{B}}^T\boldsymbol{X}^T\boldsymbol{X}\hat{\boldsymbol{B}} = \boldsymbol{V}\boldsymbol{D}\boldsymbol{V}^T$, where $\boldsymbol{V} \in \mathbb{O}^{p\times r}$ (i.e., a $p \times r$ orthogonal matrix) and set

$$\boldsymbol{Z} \triangleq \boldsymbol{X}\hat{\boldsymbol{B}}\boldsymbol{V} \in \mathbb{R}^{n\times r}. \tag{4}$$

We refer to it as the *Type-II extraction* or *post-decorrelation*. (Again, QR decomposition can be used to save some computational cost.) Each z-predictor is represented as a linear combination of the columns of $\boldsymbol{X}$, and these r newly obtained predictors are uncorrelated with each other (i.e., $\boldsymbol{Z}^T\boldsymbol{Z}$ is diagonal). Type-I and Type-II extractions are not equivalent in general, but coincide for the plain RRR/PCA estimate.

Nonasymptotic oracle inequalities can be proved for the globally optimal solutions of SEL-RRR. We use the group ℓ_0/ℓ_1 as examples to show the error rate. Define $\mathcal{J}(\boldsymbol{B}) = \{j : \boldsymbol{b}_j \neq \boldsymbol{0}\}$, $J(\boldsymbol{B}) = |\mathcal{J}(\boldsymbol{B})| = \|\boldsymbol{B}\|_{2,0}$, and $\lesssim$ denotes an inequality that holds up to a multiplicative numerical constant.

Theorem 1. *Assume $\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B}^* + \boldsymbol{E}$ with all entries of $\boldsymbol{E}$ iid $\sim \mathcal{N}(0, \sigma^2)$. Let $\hat{\boldsymbol{B}} = \arg\min_{\boldsymbol{B}:rank(\boldsymbol{B})\leq r} \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/2 + \lambda^2\|\boldsymbol{B}\|_{2,0}/2$. Then, with $\lambda = A\sigma\sqrt{r + \log p}$, where A is a constant large enough, the following oracle inequal-*

ity holds for any $\boldsymbol{B} \in \mathbb{R}^{p\times m}$ *with* $rank(\boldsymbol{B}) \le r$*:*

$$\mathbb{E}\left[\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2\right] \lesssim \|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + \lambda^2 J(\boldsymbol{B}) + (m-r)r\sigma^2 + \sigma^2. \tag{5}$$

In the ℓ_1 *case where* $\hat{\boldsymbol{B}} = \arg\min_{\boldsymbol{B}:rank(\boldsymbol{B})\le r} \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/(2\|\boldsymbol{X}\|_2) + \lambda\|\boldsymbol{B}\|_{2,1}$*, the same oracle inequality holds for any* $\boldsymbol{B} \in \mathbb{R}^{p\times m}$ *with* $rank(\boldsymbol{B}) \le r$*, provided that* $\boldsymbol{\Sigma} = \boldsymbol{X}^T\boldsymbol{X}/\|\boldsymbol{X}\|_2^2$ *satisfies Assumption* $\mathfrak{A}(\mathcal{J}(\boldsymbol{B}), \delta_{\mathcal{J}(\boldsymbol{B})}, \vartheta)$ *in Section 6.1 for some positive constants* $\delta_{\mathcal{J}(\boldsymbol{B})}$ *and* ϑ*.*

The proof is based on combined computational and statistical analysis, where Theorem 2 plays an important role. Although the rate is exemplified by the ℓ_0 and ℓ_1 penalties, our proof scheme applies much more generally: the ℓ_0 result can be extended to any $P(\theta;\lambda)$ (with λ as the threshold) that is bounded below by $P_H(\theta;\lambda) = (-\theta^2/2 + \lambda|\theta|)1_{|\theta|<\lambda} + (\lambda^2/2)1_{|\theta|\ge\lambda}$ and satisfies $P(\theta;\lambda) \le C\lambda^2$, such as the capped ℓ_1 (Zhang, 2010b) and SCAD (Fan and Li, 2001); the ℓ_1 result can be extended to any sub-additive penalty sandwiched by P_H and P_1, such as the ℓ_{p} penalties $P(\theta;\lambda) = (2-2\mathrm{p})^{1-\mathrm{p}}(2-\mathrm{p})^{\mathrm{p}-2}\lambda^{2-\mathrm{p}}|\theta|^{\mathrm{p}}$ $(0 < \mathrm{p} < 1)$. In addition, a high-probability form result with the same error bound can be obtained (without the last additive term σ^2 in (5)). See Section 6.1 for proof details.

In the group ℓ_1 case, our theorem suggests that the rate of λ at $\sigma\sqrt{r\log p}$ (Bunea et al., 2012, Theorem 3) is merely suboptimal, so is their error rate $mr^* + J^*r^*\log p$. In addition, we do not require the largest eigenvalue of $\boldsymbol{X}^T\boldsymbol{X}/n$ to be bounded.

Indeed, when $r = r^*$, setting $\boldsymbol{B} = \boldsymbol{B}^*$ such that the bias term $\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2$ vanishes, we obtain an error rate as σ^2 times $(J^* + m - r^*)r^* + J^*\log p$, where r^* and J^* are the rank and row support size of $\boldsymbol{B}^*$, respectively. (The existence of the bias term can make our conclusion applicable to *approximately* sparse signals.) When r^* is unknown, applying the PIC in Section 4 to choose the value of r (which only requires a discrete and finite search grid), we can attain the same error rate. Obviously, our result is uniformly better than $mr^* + J^*r^*\log p$ (which is suboptimal when r^* is comparable to J^*). It beats the error rate $J^*m + J^*\log p$ by group variable selection (Lounici et al., 2011) as well, and the performance gain is dramatic for low rank systems.

For pure rank reduction, (5) adds an additional p to the optimal rate $mr + qr$ (Bunea et al., 2011; Candès and Plan, 2011), probably because only large values of A are considered in the theorem. In practical data analysis, however, there is no performance loss, because when $\lambda = 0$ the criteria in Theorem 1 degenerate to the RRR.

Even sharper error rates can be achieved with a finer tuning of λ (and r) which however does not give a universal choice for λ; see Section 4.

2.2 Computational algorithms

The previous subsection reveals the statistical benefits of simultaneous rank reduction and variable selection. On the other hand, even with a convex P, (2) is not a convex optimization problem due to the discrete and nonconvex rank constraint. We address the computational issues in this subsection and develop fast, but simple-to-implement algorithms to ensure the practicality of the proposed methodologies in real-life large data applications. We start with thresholding rules to tackle the nonconvexity.

Definition 2.1 (Threshold function)**.** *A threshold function is a real valued function* $\Theta(s;\lambda)$ *defined for* $-\infty < s < \infty$ *with* λ *as the parameter (*$0 \le \lambda < \infty$*) such that*

1) $\Theta(-s;\lambda) = -\Theta(s;\lambda)$,
2) $\Theta(s;\lambda) \le \Theta(s';\lambda)$ *for* $s \le s'$,
3) $\lim_{s\to\infty} \Theta(s;\lambda) = \infty$, *and*
4) $0 \le \Theta(s;\lambda) \le s$ *for* $0 \le s < \infty$.

In words, $\Theta(\cdot;\lambda)$ is an odd monotone unbounded shrinkage rule for s, at any λ. We define $\vec{\Theta}$ as a multivariate version of the thresholding operator Θ. For any vector $\boldsymbol{a} \in \mathbb{R}^r$, $\vec{\Theta}(\boldsymbol{a};\lambda) := \boldsymbol{a}\Theta(\|\boldsymbol{a}\|_2;\lambda)/\|\boldsymbol{a}\|_2$ for $\boldsymbol{a} \neq \boldsymbol{0}$ and $\boldsymbol{0}$ otherwise. For any matrix $\boldsymbol{A} \in \mathbb{R}^{p\times r}$ with $\boldsymbol{A} = [\boldsymbol{a}_1 \ \cdots \ \boldsymbol{a}_p]^T$, $\vec{\Theta}(\boldsymbol{A};\lambda) := [\vec{\Theta}(\boldsymbol{a}_1;\lambda) \ \cdots \ \vec{\Theta}(\boldsymbol{a}_p;\lambda)]^T$.

The trick to tackle the rank constraint (2) is to write $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{V}^T$ with $\boldsymbol{V}$ being

orthogonal, i.e., $\boldsymbol{V} \in \mathbb{O}^{m\times r}$. The computation is described in Algorithm 1.

Algorithm 1: SEL-RRR/PCA Algorithm

Data: $1 \le r \le p$, $\lambda \ge 0$, $\boldsymbol{S}^{(0)} \in \mathbb{R}^{p\times r}$, Θ: a thresholding rule, M_{inner}: maximum number of inner iterations, M_{outer}: maximum number of outer iterations

(1) $t \leftarrow 0$, $K \leftarrow \|\boldsymbol{X}\|_2^2$ (with $\|\cdot\|_2$ denoting the spectral norm);

while $\|\boldsymbol{B}^{(t)} - \boldsymbol{B}^{(t-1)}\|$ *(if existing) is not small enough and* $t < M_{outer}$ **do**

(2) $t \leftarrow t+1$;

(3) Let $\boldsymbol{W} \leftarrow \boldsymbol{Y}^T\boldsymbol{X}\boldsymbol{S}^{(t)}$ and perform SVD: $\boldsymbol{W} = \boldsymbol{U}_w\boldsymbol{D}_w\boldsymbol{V}_w^T$ with $\boldsymbol{D}_w$ of size $r \times r$;

(4) $\boldsymbol{V}^{(t)} \leftarrow \boldsymbol{U}_w\boldsymbol{V}_w^T$;

(5) Perform the inner iterations

(5.1) $t' \leftarrow 0$, $\tilde{\boldsymbol{S}}^{(0)} \leftarrow \boldsymbol{S}^{(t-1)}$;

while $\|\tilde{\boldsymbol{S}}^{(t')} - \tilde{\boldsymbol{S}}^{(t'-1)}\|$ *(if existing) is not small enough and* $t' < M_{inner}$ **do**

(5.2) $t' \leftarrow t'+1$;

(5.3) $\tilde{\boldsymbol{S}}^{(t')} \leftarrow \vec{\Theta}\left(\frac{1}{K}\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}^{(t-1)} + (\boldsymbol{I} - \frac{1}{K}\boldsymbol{X}^T\boldsymbol{X})\tilde{\boldsymbol{S}}^{(t'-1)}; \lambda\right)$;

end

(5.4) $\boldsymbol{S}^{(t)} \leftarrow \tilde{\boldsymbol{S}}^{(t')}$;

(6) $\boldsymbol{B}^{(t)} \leftarrow \boldsymbol{S}^{(t)}(\boldsymbol{V}^{(t)})^T$;

end

deliver $\hat{\boldsymbol{B}} = \boldsymbol{B}^{(t)}$ and $\hat{\boldsymbol{V}} = \boldsymbol{V}^{(t)}$.

The algorithm is simple to implement and has low computational complexity. In addition to the matrix multiplication and thresholding operations, Step (3) performs an SVD, but $\boldsymbol{W}$ has only r columns, and the rank values of practical interest are usually small. In addition, for the SEL-PCA problem (3) which has an identity design, the algorithm greatly simplifies. Indeed, because $K = 1$ and $M_{inner} = 1$, the computation proceeds by

$$\begin{cases} \boldsymbol{Y}\boldsymbol{S}^{(t-1)} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T \text{ (SVD)}, \boldsymbol{V}^{(t)} \leftarrow \boldsymbol{U}\boldsymbol{V}^T \\ \boldsymbol{S}^{(t)} \leftarrow \vec{\Theta}(\boldsymbol{Y}^T\boldsymbol{V}^{(t)}; \lambda). \end{cases} \tag{6}$$

Next we show the theoretical guarantee of global convergence of Algorithm 1. Given $K > 0$, define the objective function as

$$F(\boldsymbol{B}; \lambda) = \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/(2K) + \sum_j P(\|\boldsymbol{b}_j\|_2; \lambda). \tag{7}$$

With a little abuse of notation, we also write $F(\boldsymbol{s}_1, \cdots, \boldsymbol{s}_p, \boldsymbol{V}; \lambda) = F(\boldsymbol{S}, \boldsymbol{V}; \lambda) = \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T\|_F^2/(2K) + \sum_j P(\|\boldsymbol{s}_j\|_2); \lambda)$ with $\boldsymbol{S} = [\boldsymbol{s}_1, \cdots, \boldsymbol{s}_p]^T$. Given $1 \le$

$r \le p$, $\lambda \ge 0$, $K \ge \|\boldsymbol{X}\|_2^2$, and an arbitrary starting point $\boldsymbol{S}^{(0)} \in \mathbb{R}^{p\times r}$, let the sequence of estimates generated by Algorithm 1 be $(\boldsymbol{S}^{(t)}, \boldsymbol{V}^{(t)}, \boldsymbol{B}^{(t)})$, with $\boldsymbol{S}^{(t)} = [\boldsymbol{s}_1^{(t)}, \cdots, \boldsymbol{s}_p^{(t)}]^T$, $t = 1, 2, \cdots$.

Theorem 2. *Suppose $\Theta(\cdot;\lambda)$ in Algorithm 1 is a thresholding rule (cf. Definition 2.1). Let the penalty P in the objective function satisfy*

$$P(\theta;\lambda) - P(0;\lambda) = \int_0^{|\theta|} (\sup\{s : \Theta(s;\lambda) \le u\} - u)\,\mathrm{d}u + \mathrm{q}(\theta;\lambda), \tag{8}$$

with $\mathrm{q}(\theta,\lambda)$ nonnegative and $\mathrm{q}(\Theta(s;\lambda);\lambda) = 0, \forall s \in \mathbb{R}$. Suppose the continuity assumption in Section 6.2 holds. (i) If $K \ge \|\boldsymbol{X}\|_2^2$, $F(\boldsymbol{B}^{(t)};\lambda)$ is decreasing and satisfies $rank(\boldsymbol{B}^{(t)}) \le r$ and $F(\boldsymbol{B}^{(t)}) - F(\boldsymbol{B}^{(t+1)}) \ge (1 - \|\boldsymbol{X}\|_2^2/K)\|\boldsymbol{S}^{(t)} - \boldsymbol{S}^{(t+1)}\|_F^2/2$. (ii) Furthermore, if $K > \|\boldsymbol{X}\|_2^2$, then any accumulation point of $(\boldsymbol{s}_1^{(t)}, \cdots, \boldsymbol{s}_p^{(t)}, \boldsymbol{V}^{(t)})$ is a coordinatewise minimum point of F and the function value converges monotonically to $F(\boldsymbol{s}_1^*, \cdots, \boldsymbol{s}_p^*, \boldsymbol{V}^*)$ for some coordinatewise minimum point $(\boldsymbol{s}_1^*, \cdots, \boldsymbol{s}_p^*, \boldsymbol{V}^*)$.

From the theorem, when $\mathrm{q}(\cdot;\lambda) \equiv 0$, we have a stationary point guarantee. A more careful study following She (2012) shows that the factor $(1 - \|\boldsymbol{X}\|_2^2/K)$ can be strengthened to $1 + \operatorname{ess\,inf} \mathrm{d}\Theta^{-1}(u;\lambda)/\,\mathrm{d}u - \|\boldsymbol{X}\|_2^2/K$, where $\Theta^{-1}(u;\lambda) := \sup\{s : \Theta(s;\lambda) \le u\}$ and $\operatorname{ess\,inf}$ is the essential infimum. For example, when the soft-thresholding or ℓ_1-penalty is applied, we may adopt a smaller $K = \|\boldsymbol{X}\|_2^2/2$ to speed convergence.

The framework covers a wide family of penalties through the coupling equation (8) between Θ and P. The function $\mathrm{q}(\cdot)$ is often zero, but can be nonzero when Θ has discontinuities. In fact, we can use a nontrivial $\mathrm{q}(\cdot)$ to attain the exact group ℓ_0 penalty. The theorem implies that multiple (infinitely many, as a matter of fact) penalties can result in the same solution, which justifies our thresholding launching point. Employing various thresholding rules, one can reach commonly used penalties, including group versions of the ℓ_1, ℓ_0, SCAD (Fan and Li, 2001), MCP (Zhang, 2010a), capped ℓ_1 (Zhang, 2010b), $\ell_{\mathrm{p}}(0 < \mathrm{p} < 1)$, elastic net (Zou and Hastie, 2005), Berhu (Owen, 2007), etc., as demonstrated in She (2012). For example, for $P(\theta;\lambda) = (2-2\mathrm{p})^{1-\mathrm{p}}(2-\mathrm{p})^{\mathrm{p}-2}\lambda^{2-\mathrm{p}}|\theta|^{\mathrm{p}}$ $(0 < \mathrm{p} < 1)$, $\Theta(s;\lambda)$ is the largest zero crossing (in magnitude) of $g(s) = s - y + \mathrm{p}(2-2\mathrm{p})^{1-\mathrm{p}}(2-\mathrm{p})^{\mathrm{p}-2}\lambda^{2-\mathrm{p}}|s|^{\mathrm{p}-1}$ if $|s| > \lambda$, and 0 if $|s| \le \lambda$. In particular, we introduce the *hard-ridge* thresholding rule that fuses the hard-thresholding and the ridge-thresholding

$$\Theta_{HR}(s;\lambda,\eta) = \begin{cases} 0, & \text{if } |s| < \lambda \\ \frac{s}{1+\eta}, & \text{if } |s| \ge \lambda. \end{cases} \tag{9}$$

Setting $\mathrm{q}(\theta;\lambda,\eta) = \frac{(1+\eta)(\lambda-|\theta|)^2}{2}$ if $0 < |\theta| < \lambda$ and 0 if $\theta = 0$ or $|\theta| > \lambda$, we obtain the '$\ell_0 + \ell_2$' penalty: $P(\theta) = \frac{1}{2}\eta\theta^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}1_{\theta\neq 0}$, which corresponds to

$$\frac{1}{2}\eta\|\boldsymbol{B}\|_F^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}\|\boldsymbol{B}\|_{2,0} \tag{10}$$

in (2). This hybrid 'hard-ridge' penalty may be of interest in statistical learning tasks that have joint concerns of accuracy and parsimony: the group ℓ_0 portion enforces parsimonious selection, while the ridge (Frobenius) portion shrinks $\boldsymbol{B}$ to compensate for large noise and decorrelates the input variables in large-p applications. Section 4 makes use of a constraint type of this regularization.

The initial point of $\boldsymbol{S}$ can affect the final solution due to nonconvexity. In general, we construct $\boldsymbol{S}^{(0)}$ as $(\boldsymbol{X}^T\boldsymbol{X})^+\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}_r$, based on the RRR estimate $\hat{\boldsymbol{B}} = (\boldsymbol{X}^T\boldsymbol{X})^+\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}_r\boldsymbol{V}_r^T$, where $\boldsymbol{V}_r$ is formed by the first r eigenvectors of $\boldsymbol{Y}^T\mathbf{P}_{\boldsymbol{X}}\boldsymbol{Y}$. (In the PCA setting (3), this corresponds to using the first r left singular vectors of $\boldsymbol{X}$ for $\boldsymbol{V}^{(1)}$.) Of course, other initializations are possible—for example, one can use the leading PC directions of a submatrix (pre-screened according to column norms) to initialize $\boldsymbol{S}^{(0)}$, see Johnstone and Lu (2009) and Ma (2013).

Nicely, our SV formulation can also characterize sparse RRR/PCA. Concretely, with $\vec{\Theta}$ replaced by Θ applied componentwise in Step (5.3), we can prove a similar result to Theorem 2 (details omitted), which shows that the corresponding computational algorithm solves a *sparse reduced rank regression* (**SPA-RRR**) problem

$$\textbf{SPA-RRR:}\quad \min_{\boldsymbol{S}=[s_{j,k}]\in\mathbb{R}^{p\times r},\boldsymbol{V}\in\mathbb{O}^{m\times r}} \|\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T\|_F^2/(2K) + \sum_{j,k} P(|s_{j,k}|;\lambda^e) \tag{11}$$

for any P satisfying (8). λ^e can be replaced by λ_k^e or $\lambda_{j,k}^e$. Letting $\boldsymbol{S} = [\tilde{\boldsymbol{s}}_1,\cdots,\tilde{\boldsymbol{s}}_r]$ and $\boldsymbol{V} = [\boldsymbol{v}_1,\cdots,\boldsymbol{v}_r]$, we have

$$\boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T = \boldsymbol{X}\tilde{\boldsymbol{s}}_1\boldsymbol{v}_1^T + \cdots + \boldsymbol{X}\tilde{\boldsymbol{s}}_r\boldsymbol{v}_r^T.$$

Each term on the right side is of rank 1, and involves only a subset of predictors chosen by the nonzero entries in $\tilde{\boldsymbol{s}}_k$ ($1 \le k \le r$).

In the unsupervised setting with only $\boldsymbol{X}$ available, a special case of (11) gives a criterion for rank-r sparse PCA:

$$\textbf{SPA-PCA:}\quad \min_{\boldsymbol{S}\in\mathbb{R}^{p\times r},\boldsymbol{V}\in\mathbb{O}^{p\times r}} \|\boldsymbol{X}-\boldsymbol{V}\boldsymbol{S}^T\|_F^2/(2K) + \sum_{j,k} P(|s_{j,k}|;\lambda^e). \tag{12}$$

Here, the iteration has a simple form in implementation (cf. (6), with $\vec{\Theta}$ replaced by Θ). One can obtain a multivariate sparse PCA solution and avoid the *ad hoc* design of deflation (Mackey, 2008). Of course, a sequential algorithm can be developed as well. Indeed, when $r = 1$, the objective functions of SEL-PCA and SPA-PCA coincide. Since the $\boldsymbol{W}$ and $\boldsymbol{V}^{(t)}$ in Steps (3)-(4) are now vectors, there is no need to perform the SVD in the algorithm: $\boldsymbol{V}^{(t)}$ is simply $\boldsymbol{W}/\|\boldsymbol{W}\|_2$.

Some Existing Methods and Models for Sparse PCA. SEL-RRR/PCA can construct multiple factors simultaneously from a reduced subset of input variables. The same idea of pursing joint sparsity to guarantee selectability is used in the sparse PCA by Cai et al. (2013), where an adaptive multi-step procedure is proposed under the spiked covariance model assumption. Our computational algorithm has an associated optimization criterion and allows for any thresholding function. SEL-PCA is close in spirit to the algorithm design in the seminal work by Johnstone and Lu (2009) which performs simple thresholding according to the *variance* of each variable, followed by a reduced PCA (and a complementary filtering).

Although this paper focuses on the SEL-RRR/PCA, there is a large body of literature relevant to its SPA variant. Our rank-r SPA-PCA shares some similarities with the ITSPCA for principal subspace estimation (Ma, 2013). ITSPCA does not have a formulated optimization criterion as (12) and does not have a theoretical guarantee of computational convergence. Interestingly, Ma (2013) showed that one can terminate the algorithm after $\mathcal{O}(\log n)$ steps and obtain small subspace error with high probability.

There are much more works in the literature to solve the rank-1 sparse PCA problem. An effective trick used in Shen and Huang (2008), Witten et al. (2009), and Journée et al. (2010) is to introduce $\boldsymbol{v}' \in \mathbb{R}^p$ and replace the first term $\boldsymbol{v}^T(\boldsymbol{X}^T\boldsymbol{X})\boldsymbol{v}$ by $\boldsymbol{v}'^T\boldsymbol{X}\boldsymbol{v}$ with $\boldsymbol{v}'$ free of regularization, or $(\boldsymbol{v}'^T\boldsymbol{X}\boldsymbol{v})^2$ for $\boldsymbol{v}' : \boldsymbol{v}'^T\boldsymbol{v}' = 1$. Then alternative optimization can be applied. But as mentioned previously, all are subject to the design of deflation, and easily result in excessive dimensions when r is not too low.

We finally comment on a regularized 'self-regression' form for sparse PCA (e.g., Zou et al. (2006)), which is attained by setting $\boldsymbol{Y} = \boldsymbol{X}$ in (11). This may however result in serious support ambiguity. Assume a noise-free sparse PCA model with $\boldsymbol{X}_{(I^*)^c} = \boldsymbol{0}$ and $\text{rank}(\boldsymbol{X}_{I^*}) = r$. Then, any variable index set $I \subset [p] =: \{1, \cdots, p\}$ satisfying $\text{rank}(\boldsymbol{X}_I) = r$ leads to $\boldsymbol{X} = \boldsymbol{X}\boldsymbol{B}$ for some $\boldsymbol{B}$ satisfying $\boldsymbol{B}_{I^c} = \boldsymbol{0}$ and $\text{rank}(\boldsymbol{B}_I) \leq r$. (In fact, in the SVD $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T$ where $\boldsymbol{D}$ is an $r \times r$ diagonal matrix, $\boldsymbol{V}_I$ must have full rank r and thus $\boldsymbol{V}_{I^c} = \boldsymbol{V}_I\boldsymbol{A}$ for some $\boldsymbol{A}$, or $[\boldsymbol{X}_I, \boldsymbol{X}_{I^c}] = \boldsymbol{X}_I\boldsymbol{V}_I\boldsymbol{V}_I^T[\boldsymbol{I},\ \boldsymbol{A}] = [\boldsymbol{X}_I, \boldsymbol{X}_{I^c}][\boldsymbol{B}_I^T, \boldsymbol{0}]^T$ with $\boldsymbol{B}_I =$

$\boldsymbol{V}_I\boldsymbol{V}_I^T[\boldsymbol{I},\ \boldsymbol{A}]$.) Therefore, the row support of $\boldsymbol{B}$ in the self-regression $\boldsymbol{X} = \boldsymbol{X}\boldsymbol{B}$ can even be independent of the I^* used in the model specification.

3 Model comparison by a predictive information criterion

In selectable reduced rank analysis, we need to tune two parameters for variable selection and rank reduction purposes, respectively. We design an information criterion from a predictive learning perspective—the best model should minimize the prediction error. As will be shown later, such a principle can avoid model ambiguity issues and signal-to-noise requirement. Indeed, unlike those for the purpose of consistent rank/variable selection, our criterion will put no restrictions or regularity conditions on $\boldsymbol{X}$ or $\boldsymbol{B}$. This shares similarity with Foster and George (1994), where the well-known RIC is developed in the regression setup (assuming however $n - p$ large). Our nonasymptotic study in this section could be viewed as an extension of RIC to vector models with joint concerns of rank reduction and variable selection. We refer to it as the predictive information criterion (**PIC**).

We assume

$$\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}^* = \boldsymbol{E}, \tag{13}$$

where $\boldsymbol{Y} \in \mathbb{R}^{n\times m}$, $\boldsymbol{X} \in \mathbb{R}^{n\times p}$, and vec $(\boldsymbol{E})$ is *sub-Gaussian* with mean zero and scale bounded by σ; see Definition 6.1 in Section 6.3. Note that column and/or row dependencies in $\boldsymbol{E}$ are allowed.

All information criteria must append a model complexity penalty to the fitting error term. To motivate the design of PIC, we will first prove a nonasymptotic oracle inequality. For notational simplicity, use $r(\cdot)$ to denote rank$(\cdot)$ and efine

$$P_o(\boldsymbol{B}) = \sigma^2\{[q \wedge J(\boldsymbol{B}) + m - r(\boldsymbol{B})]r(\boldsymbol{B}) + J(\boldsymbol{B})\log(ep/J(\boldsymbol{B})\}, \tag{14}$$

where $q := \text{rank}(\boldsymbol{X})$.

Theorem 3. *Under model assumption* (13) *with the sub-Gaussian noise contamination, let* $\hat{\boldsymbol{B}} \in \arg\min_{\boldsymbol{B}} \frac{1}{2}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2 + AP_o(\boldsymbol{B})$*, where* A *is a constant. Then for all sufficiently large values of* A*,* $\hat{\boldsymbol{B}}$ *satisfies a nonasymptotic oracle inequality for any* $\boldsymbol{B} \in \mathbb{R}^{p\times m}$

$$\mathbb{E}\left[\max(\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2, P_o(\hat{\boldsymbol{B}}))\right] \lesssim \inf_{\boldsymbol{B}}\{\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + P_o(\boldsymbol{B})\} + \sigma^2. \tag{15}$$

The proof in Section 6.3 is based on a decomposition trick and a covering argument. The first term in P_o, $[q \wedge J(\boldsymbol{B}) + m - r(\boldsymbol{B})]r(\boldsymbol{B})$, gives the total number of free parameters (or degrees of freedom, **DF**) in $\boldsymbol{XB}$, and the second term $J(\boldsymbol{B}) \log(ep/J(\boldsymbol{B})$ characterizes risk inflation (denoted by **IF**). We emphasize that this nonasymptotic oracle result is under no restrictions on the design matrix or the true coefficient matrix and applies to *any* reference signal $\boldsymbol{B}$.

Theorem 3 is significantly stronger than Theorem 1 of Bunea et al. (2012) which relies on a multiplicative penalty $r(\boldsymbol{B})J(\boldsymbol{B}) \log(ep/J(\boldsymbol{B})) + mr(\boldsymbol{B})$ that is larger than P_o. (There exists a derivation error in their proof of Theorem 1—see Section 6.3.)

Accordingly, we get a nonasymptotic PIC for model selection. Indeed, under the model assumption (13), for any collection of (random) non-zero matrices $\boldsymbol{B}_1$, $\cdots$, $\boldsymbol{B}_l$, $\cdots$, choose the optimal one $\boldsymbol{B}_o$ by minimizing the following criterion over all $\boldsymbol{B}_l$

$$\textbf{PIC:} \qquad \|\boldsymbol{Y} - \boldsymbol{XB}\|_F^2 + AP_o(\boldsymbol{B}), \tag{16}$$

with the penalty P_o defined in (14) and the constant A appropriately large. Then by the same reasoning as in the proof of Theorem 3 (details omitted), $\boldsymbol{B}_o$ satisfies

$$\mathbb{E}\left[\|\boldsymbol{XB}_o - \boldsymbol{XB}^*\|_F^2\right] \le C \inf_l \mathbb{E}[\|\boldsymbol{XB}_l - \boldsymbol{XB}^*\|_F^2 + P_o(\boldsymbol{B}_l)].$$

The IF with the multiplicative logarithm factor is purely due to variable selection. (In fact, for sole rank reduction, we can similarly derive a PIC in which the selection inflation disappears.) For the univariate response model with $m = 1$ and $r = 1$, P_o gives a finer version of the RIC which is at the rate of $\mathcal{O}(\log p)J(\boldsymbol{B})\sigma^2$. Perhaps interestingly, for multivariate models, the IF term is additive to the DF term, and does not involve rank information at all. This seems to be different than the information criteria we are familiar with in the area of variable selection. For example, EBIC (Chen and Chen, 2008) for high-dimensional variable selection has a multiplicative $\log p$ factor on the whole DF of the model.

The scale parameter σ is unknown in the above selection criterion. One can plug in a good estimate $\hat{\sigma}$ for the unknown σ—for example, $\hat{\sigma}^2 = \text{med}(\|\boldsymbol{x}_j\|_2^2/n)$ (Johnstone and Lu, 2009) can be used in the selectable/sparse PCA.

Furthermore, we propose a *scale-free* PIC (**SF-PIC**) for parsimonious model tuning, which bypasses the estimation of the scale parameter. For simplicity, we assume $\boldsymbol{E}$ has iid Gaussian entries $\sim \mathcal{N}(0, \sigma^2)$. Still, no coherence condition is required for the design in the following theorem.

Theorem 4. *Suppose the true model is parsimonious in the sense that* $P_o(\boldsymbol{B}^*) < mn\sigma^2/A_0$ *for some constant* $A_0 > 0$*. Consider a criterion defined by*

$$\textbf{SF-PIC:} \qquad \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/\{mn - AP_o(\boldsymbol{B})/\sigma^2\}, \tag{17}$$

where A *is a constant satisfying* $0 < A < A_0$*. For sufficiently large enough values of* A_0 *and* A*, any* $\hat{\boldsymbol{B}}$ *that minimizes* (17) *achieves the prediction error rate:* $\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \lesssim P_o(\boldsymbol{B}^*)$*, with probability at least* $1 - Cp^{-c} - C'\exp(-c'mn)$ *for some constants* $C, C', c, c' > 0$*.*

Remark. Our analysis only reveals the complexity rate in PIC (or SF-PIC), i.e., A_1**DF** $+ A_2$**IF**, with A_1, A_2 to be determined. However, because A_1 and A_2 are constants rather than tuning parameters, one can figure out their values by computer experiments. Our experience shows that for PIC where σ^2 is known or can be well estimated, the choice of $A_1 = 2.4$, $A_2 = 1.8$ works well, and for SF-PIC we recommend $A_1 = 2$, $A_2 = 1.8$.

4 Rank constrained variable screening in ultrahigh dimensional computation

Recently, big data computing attracts a lot of attention from statisticians. In this section, we propose a novel means for rank-constrained variable screening which can remarkably reduce the computational cost for applications in ultrahigh dimensions. The key is to change the group selection penalty to a group ℓ_0 constraint and add in progressive squeezing operations.

We introduce a *quantile thresholding rule* $\Theta^{\#}(\cdot; \mathfrak{q}, \eta)$ as a variant of the hard-ridge thresholding (9). Given $1 \le \mathfrak{q} \le p$ and $\eta \ge 0$, $\Theta^{\#}(\boldsymbol{a}; \mathfrak{q}, \eta) : \mathbb{R}^p \to \mathbb{R}^p$ is defined for any $\boldsymbol{a} \in \mathbb{R}^p$ such that the $\mathfrak{q}$ largest components of $\boldsymbol{a}$ (in absolute value) are shrunk by a factor of $(1 + \eta)$ and the remaining components are all set to be zero. In the case of ties, a random tie breaking rule is used. A matrix version $\vec{\Theta}^{\#}$ to be used in our problem is defined as

$$\vec{\Theta}^{\#}(\boldsymbol{A}; \mathfrak{q}, \eta) = \text{diag}\{\Theta^{\#}(\boldsymbol{g}(\boldsymbol{A}); \mathfrak{q}, \eta)\}\boldsymbol{A}^{\circ},$$

for any $\boldsymbol{A} = [\boldsymbol{a}_1, \cdots, \boldsymbol{a}_p]^T \in \mathbb{R}^{p\times r}$, where $\boldsymbol{g}(\boldsymbol{A}) := [\|\boldsymbol{a}_j\|_2]_{p\times 1}$ and $\boldsymbol{A}^{\circ} = (\text{diag}\{\boldsymbol{g}(\boldsymbol{A})\})^{+}\boldsymbol{A}$ with $^{+}$ standing for the Moore-Penrose pseudoinverse. Algo-

rithm 2 is described as follows.

Algorithm 2: SEL-RRR/PCA Screening Algorithm

Data: $1 \leq r \leq p$, $\eta \geq 0$, $1 \leq \mathfrak{q} \leq p$, $\boldsymbol{S}^{(0)} \in \mathbb{R}^{p\times r}$, M_{inner}/M_{outer}: maximum number of inner/outer iterations

(1) $t \leftarrow 0$, $K \leftarrow \|\boldsymbol{X}\|_2^2$;

while $\|\boldsymbol{B}^{(t)} - \boldsymbol{B}^{(t-1)}\|$ *(if existing) is not small enough and* $t < M_{outer}$ **do**

(2) $t \leftarrow t+1$;

(3) Let $\boldsymbol{W} \leftarrow \boldsymbol{Y}^T\boldsymbol{X}\boldsymbol{S}^{(t)}$ and perform SVD: $\boldsymbol{W} = \boldsymbol{U}_w\boldsymbol{D}_w\boldsymbol{V}_w^T$ with $\boldsymbol{D}_w$ of size $r \times r$;

(4) $\boldsymbol{V}^{(t)} \leftarrow \boldsymbol{U}_w\boldsymbol{V}_w^T$;

(5) Perform the inner iterations

(5.1) $t' \leftarrow 0$, $\tilde{\boldsymbol{S}}^{(0)} \leftarrow \boldsymbol{S}^{(t-1)}$;

while $\|\tilde{\boldsymbol{S}}^{(t')} - \tilde{\boldsymbol{S}}^{(t'-1)}\|$ *(if existing) is not small enough and* $t' < M_{inner}$ **do**

(5.2) $t' \leftarrow t'+1$;

(5.3) $\tilde{\boldsymbol{S}}^{(t')} \leftarrow \vec{\Theta}^{\#}\left(\frac{1}{K}\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}^{(t-1)} + (\boldsymbol{I} - \frac{1}{K}\boldsymbol{X}^T\boldsymbol{X})\tilde{\boldsymbol{S}}^{(t'-1)}; \mathfrak{q}, \eta\right)$;

end

(5.4) $\boldsymbol{S}^{(t)} \leftarrow \tilde{\boldsymbol{S}}^{(t')}$;

(6) $\boldsymbol{B}^{(t)} \leftarrow \boldsymbol{S}^{(t)}(\boldsymbol{V}^{(t)})^T$;

end

deliver $\hat{\boldsymbol{B}} = \boldsymbol{B}^{(t)}$ and the remaining dimensions determined by $\{j : \boldsymbol{g}(\boldsymbol{B}^{(t)})[j] \neq 0\}$.

Given $1 \leq r \leq p$, $1 \leq \mathfrak{q} \leq p$, $\eta \geq 0$, and an arbitrary starting point $\boldsymbol{S}^{(0)} \in \mathbb{R}^{p\times r}$, denote by $(\boldsymbol{S}^{(t)}, \boldsymbol{V}^{(t)}, \boldsymbol{B}^{(t)})$ $(t = 1, 2, \cdots)$ the sequence of iterates generated by Algorithm 2. Let $F(\boldsymbol{B}) = \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/(2K) + \eta\|\boldsymbol{B}\|_F^2/2$ be the objective function.

Theorem 5. *Suppose* $K \geq \|\boldsymbol{X}\|_2^2$. *Then* $F(\boldsymbol{B}^{(t)})$ *is decreasing:* $F(\boldsymbol{B}^{(t)}) - F(\boldsymbol{B}^{(t+1)}) \geq (1 - \|\boldsymbol{X}\|_2^2/K)\|\boldsymbol{S}^{(t)} - \boldsymbol{S}^{(t+1)}\|_F^2/2$, *and* $\boldsymbol{B}^{(t)}$ *obeys* $rank(\boldsymbol{B}^{(t)}) \leq r$ *and* $\|\boldsymbol{B}^{(t)}\|_{2,0} \leq \mathfrak{q}$, *for any* $t \geq 1$.

Therefore, Algorithm 2 conducts rank constrained variable screening by solving a constraint form of (3):

$$\min_{\boldsymbol{B}} \frac{1}{2K}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2 + \frac{\eta}{2}\|\boldsymbol{B}\|_F^2 \text{ s.t. } \|\boldsymbol{B}\|_{2,0} \leq \mathfrak{q}, rank(\boldsymbol{B}) \leq r. \tag{18}$$

A finite-sample oracle inequality similar to Theorem 1 can be shown following the lines of Section 6.1. This type of problems are of great interest in (ultra)high-dimensional applications because under the sparsity assumption it is not difficult for one to specify an upper bound $\mathfrak{q}$ for the number of relevant features

$\|\boldsymbol{B}^*\|_{2,0}$. Usually, $\mathfrak{q}$ satisfies $\mathfrak{q} < n$, and so after the screening process, one faces a low-dimensional problem. Running the SPA-RRR or SPA-PCA on the screened dataset significantly reduces the computation burden.

Algorithm 2 is closely connected with the SIS (Fan and Lv, 2008). Given a univariate response (denoted by $\boldsymbol{y}$) model, with $\boldsymbol{S}^{(0)} = \boldsymbol{0}$, the first iteration ranks all features based on $\boldsymbol{X}^T\boldsymbol{y}$, which amounts to the SIS. Of course, this screening is purely based on marginal statistics, and may be improper in the presence of many (correlated) predictors. Algorithm 2 iterates to lessen such greediness. The iterative quantile screening technique applies to the sole variable selection, group variable selection, or rank reduction.

We adopt an '*annealing + squeezing*' idea to further reduce computational load for big data. Let $\mathfrak{q}$ be an upper bound of the number of relevant x variables and r be an upper bound of the number of necessary underlying factors. Define a cooling schedule $Q(T)$ ($1 \le T \le \bar{T}$) with $Q(1) = p$ and $Q(\bar{T}) = \mathfrak{q}$. We propose a **progressive** quantile thresholding-based iterative screening procedure as follows.

1. Initialization: $\boldsymbol{d} \leftarrow [1\ 2\ \cdots\ p]$.
2. Given each T, run Steps (1)-(6) of Algorithm 2 with $\mathfrak{q}$ replaced by $Q(T)$, and add 'squeezing' operations afterwards:

   (7) $\tilde{\boldsymbol{d}} \leftarrow \{j : \boldsymbol{g}(\boldsymbol{S}^{(t)})[j] \neq 0\}, \boldsymbol{d} \leftarrow \boldsymbol{d}[\tilde{\boldsymbol{d}}], \boldsymbol{S}^{(t)} \leftarrow \boldsymbol{S}^{(t)}[\tilde{\boldsymbol{d}},], \boldsymbol{X} \leftarrow \boldsymbol{X}[,\tilde{\boldsymbol{d}}]$.
3. Repeat this for $T = 1, \cdots, \bar{T}$, and deliver the final remaining dimensions indexed by $\boldsymbol{d}$.

Because of the squeezing operations, the sizes of $\boldsymbol{W}$ and $\boldsymbol{S}$ keep dropping as t increases. This is particularly helpful for very high dimensional data. Empirically, we set $M_{outer} = M_{inner} = 1$, and find the sigmoidal decay cooling schedule $Q(t) = [2p/(1 + \exp(\alpha T))]$ with $\alpha = 0.01$ achieve good balance between selection and efficiency.

Similar to the discussion of Algorithm 1, if in Step (5.3) the quantile thresholding $\Theta^{\#}(\cdot; \mathfrak{q}; \eta)$ is applied after vectorizing $\frac{1}{K}\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}^{(t-1)} + (\boldsymbol{I} - \frac{1}{K}\boldsymbol{X}^T\boldsymbol{X})\tilde{\boldsymbol{S}}^{(t'-1)}$, an analogue of Theorem 5 can be proved and the algorithm solves a *constraint* type of **SPA-RRR** (or **SPA-PCA** with an identity design matrix):

$$\min_{\boldsymbol{S}\in\mathbb{R}^{p\times r}, \boldsymbol{V}\in\mathbb{O}^{m\times r}} \frac{1}{2K}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T\|_F^2 + \frac{\eta}{2}\|\boldsymbol{S}\|_F^2 \text{ s.t. } \|\operatorname{vec}(\boldsymbol{S})\|_0 \le \mathfrak{q}^e. \qquad (19)$$

We advocate a *hybrid* $\boldsymbol{S}^2$*-strategy* to combine the virtues of SEL-RRR and SPA-RRR. First, to tackle the ultrahigh dimensionality under rank constraint, run SEL-RRR (e.g., the progressive quantile version) for screening. Then, on the reduced low dimensional data, run SPA-RRR to get all factors and loading directions with individual sparsity patterns.

5 Experiments

Paper quality data. This is the first of the two datasets collected from Norwegian paper industry in Aldrin (1996). 13 measures are used to evaluate paper quality. The production of paper depends on a huge number of variables, and the data were obtained by varying three controllable variables $\boldsymbol{X}_1$, $\boldsymbol{X}_2$, $\boldsymbol{X}_3$, whose values are denoted by 1 (high level), 0, or 1 (low level). There are 15 design points, and at each design point, the paper quality was measured twice, except once for $(0, 1, 1)$, resulting in 29 observations in total. Aldrin suggests a full quadratic model with 9 predictors and 13 responses.

Previous analyses, e.g., Aldrin (1996), Izenman (2008), have revealed that the model is factor driven. One can apply RRR for *supervised* factor extraction, and obtain explanatory score variables for judging paper quality. On the other hand, it is helpful to find out the truly relevant predictors, which amounts to selectable factor extraction by SEL-RRR. We split the data into a training subset (60%) and a test subset (40%) for 100 times, and compare the predictive performance of the two methods, RRR and SEL-RRR, both tuned by SF-PIC. The median statistics are as follows. The test error of RRR is 166.8 and its rank is 2; SEL-RRR gives the same number of factors and achieves a comparable test error 167.8, using however about half of the predictors (5).

A careful examination of the data shows some interesting findings that merit further investigation. Although observation 5 and observation 6 have the same design point $(0, 0, 0)$, some of their responses show larger-than-normal discrepancies—for example, in the first response, the $(5, 6)$ between-pair difference is -1.7831, while all the other observation pairs have the differences within $[-0.4458, 0.4458]$.

From the data description (and in comparison with the second paper dataset in Aldrin (1996)), we suspect that the issue is due to the crude coding of the control variables when they are varied. A perhaps more appropriate model is driven by robust factors. That is, the factors based on $\boldsymbol{X}$, say $\boldsymbol{F} = \boldsymbol{X}\boldsymbol{A}$, may need some correction, $\boldsymbol{O}$, at few anomaly points, which amounts to $\boldsymbol{Y} = (\boldsymbol{F} + \boldsymbol{O})\boldsymbol{C} + \boldsymbol{E}$ or $\boldsymbol{Y} = \bar{\boldsymbol{X}}\bar{\boldsymbol{B}} + \boldsymbol{E}$, where $\bar{\boldsymbol{X}} = [\boldsymbol{X}\ \boldsymbol{I}]$. In this model, p is always larger than n and

Table 1: SPA-PCA and $\mathbf{S}^2$-PCA with $\mathfrak{q} = \mathfrak{q}^e$, $r = 30$ in terms of computational time (in seconds), adjusted variance (AV), and # of pixels involved.

$\mathfrak{q}^e$		7200	4800	2400	1200
Time	SPA-PCA	239	254	278	293
	$\mathbf{S}^2$-PCA	204	131	87	81
AV	SPA-PCA	64.1	51.3	33.4	20.8
	$\mathbf{S}^2$-PCA	64.0	51.1	35.9	20.5
Pixels	SPA-PCA	5742	4435	2370	1198
	$\mathbf{S}^2$-PCA	5658	4377	2187	1127

$\bar{\boldsymbol{B}}$ is desired to be row sparse. We re-run SEL-RRR using this augmented design. The median number of the selected variables is still 5, but among the predictor list only $\boldsymbol{X}_1$ and $\boldsymbol{X}_3$ are selected, and observation 6 and observation 7, at design points $(0, 0, 0)$ and $(-1, 0, 1)$, respectively, are identified as anomalies more than 50% of the time. The median rank value is lower (1) and the test error is substantially reduced to 147.5.

Face data. The Extended Yale Face Database B (Lee et al. (2005), Georghiades et al. (2001)) contains aligned and cropped face images of 38 subjects under 64 different illumination conditions (with the same frontal pose). Each image contains 192×168 pixels; we down-sampled them to $96 \times 84 (= 8064)$ to save some computational time. Given a subject ID, we extract all his/her 64 facial images to make a data matrix $\boldsymbol{Y}$ of size 64$\times$8,064.

In face recognition, PCA is widely used to extract r (say $r = 30$) basis features, referred to as *eigenfaces*, for dimension reduction and classification. Here, we aim to use selectable/sparse PCA to capture the variation and the regions of interest regarding different light source directions (in azimuth and elevation angles). We focus on one subject, with ID $= 22$ (say). Notice that the ultrahigh dimensionality $8,064$ (relative to 64) poses a great challenge in computation. Some examples of his face images are shown in the top panel of Figure 1, together with a demonstration for SEL-PCA with $r = 30$ and $\mathfrak{q} = 2400, 1200, 600, 200$ in the middle panel. Interestingly, SEL-PCA automatically identified the most informative regions, such as nose tip and forehead, that are sensitive to different illumination conditions, which shows an intuitive example of joint feature selection and extraction.

Table 1 shows the computational benefit of the advocated $\mathbf{S}^2$ strategy. In the experiment, SPA-PCA on the original dataset is compared with $\mathbf{S}^2$-PCA by set-

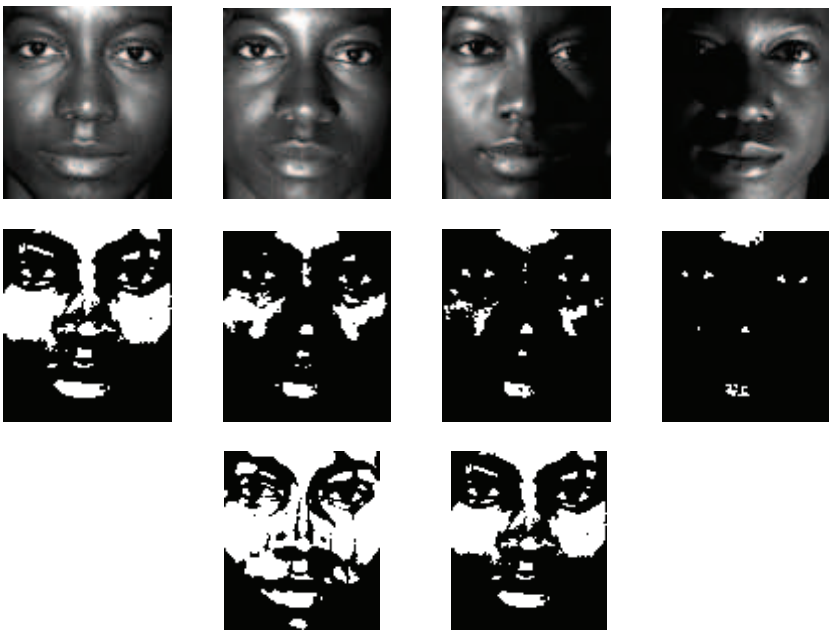

Figure 1: Top panel: Face image examples for subject 22 with image IDs given by 1, 10, 20, 50. Middle panel: Regions of interest determined by SEL-PCA (denoted by white pixels) with $r = 30$, $\mathfrak{q} = 2400, 1200, 600, 200$ (left to right). Bottom panel: Remaining pixels of SPA-PCA with $r = 30$, $\mathfrak{q}^e = 3600$ (left) and $\mathbf{S}^2$-PCA with $r = 30$, $\mathfrak{q} = 2400$, $\mathfrak{q}^e = 3600$ (right).

ting $\mathfrak{q} = \mathfrak{q}^e$ in the SEL-PCA followed by the SPA-PCA on the *reduced* dataset. (Recall that $\mathfrak{q}$ is the number of nonzero rows in $\boldsymbol{S}$ in the SEL-PCA, $\mathfrak{q}^e$ is the number of nonzero elements in $\boldsymbol{S}$ in the (non-sequential) SPA-PCA , and $\mathfrak{q}^e \leq \mathfrak{q}r$ in the $\mathbf{S}^2$-PCA.) Although $\mathfrak{q} = \mathfrak{q}^e$ is a conservative choice in screening, impressive CPU time savings are achieved in Table 1, especially when $\mathfrak{q}^e$ is not very large (relative to p). Moreover, $\mathbf{S}^2$-PCA gives essentially the same adjusted variances (Shen and Huang, 2008) but uses fewer pixels. Furthermore, an $\mathbf{S}^2$-PCA experiment with $\mathfrak{q} < \mathfrak{q}^e$ is shown in the bottom panel of Figure 1, where $\mathfrak{q}^e = 3600$ and $\mathfrak{q} = 2400$. Though having similar adjusted variances (AV $= 43\%, 40\%$, respectively), SPA-PCA used 3517 pixels in total, while $\mathbf{S}^2$-PCA reduced the number by about one third (using 2400 pixels in total). Moreover, we found that the SPA-PCA involving 2400 pixels only gave an AV of 34%.

6 Proofs

6.1 Proof of Theorem 1

Throughout the paper, we use C, c, L to denote universal constants. They are not necessarily the same at each occurrence.

Let $P_1(t;\lambda) = \lambda|t|$, $P_0(t;\lambda) = \frac{\lambda^2}{2}1_{t\neq 0}$, $P_H(t;\lambda) = (-t^2/2 + \lambda|t|)1_{|t|<\lambda} + (\lambda^2/2)1_{|t|\geq\lambda}$. Both $P_1(t;\lambda)$ and $P_0(t;\lambda)$ are bounded below by $P_H(t;\lambda)$. For convenience, $P_{2,1}(\boldsymbol{B};\lambda) = \lambda\|\boldsymbol{B}\|_{2,1} = \sum \lambda\|\boldsymbol{b}_j\|_2$; $P_{2,0}$ and $P_{2,H}$ are defined similarly. Consider a scaled SEL-RRR problem:

$$\min_{\boldsymbol{B}:rank(\boldsymbol{B})\leq r} \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2/(2\rho^2) + \sum P(\|\boldsymbol{b}_j\|_2;\lambda/\rho), \tag{20}$$

where $\rho = \|\boldsymbol{X}\|_2$ and P is P_0 or P_1. For short, we also write the penalty term as $P(\boldsymbol{B};\lambda/\rho)$ when there is no ambiguity. By definition, the estimator $\hat{\boldsymbol{B}}$ satisfies

$$\frac{1}{2}\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \leq \frac{1}{2}\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + \rho^2(P(\boldsymbol{B};\frac{\lambda}{\rho}) - P(\hat{\boldsymbol{B}};\frac{\lambda}{\rho})) + \langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle \tag{21}$$

where $\boldsymbol{\Delta} = \hat{\boldsymbol{B}} - \boldsymbol{B}$ satisfying $r(\boldsymbol{\Delta}) \leq 2r$. We use the following lemma to bound the last stochastic term.

Lemma 1. *For any given $1 \leq J \leq p, 1 \leq r \leq m \wedge J$, define $\Gamma_{J,r} = \{\boldsymbol{A} \in \mathbb{R}^{p\times m} : rank(\boldsymbol{A}) \leq r, J(\boldsymbol{A}) = J\}$. Then there exist universal constants $A_0, C, c > 0$ such that for any $a \geq \sqrt{2/A_0}$, the following event*

$$\sup_{\boldsymbol{A}\in\boldsymbol{\Gamma}_{J,r}} \{2\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{A}\rangle - \frac{1}{a}\|\boldsymbol{X}\boldsymbol{A}\|_F^2 - aA_0[\rho^2 P_{2,H}(\boldsymbol{A};\frac{\lambda^o}{\rho}) + \sigma^2 r(m-r)]\} \geq a\sigma^2 t \tag{22}$$

occurs with probability at most $C\exp(-ct)$ for any $t \geq 0$, where $\lambda^o = \sigma\sqrt{r + \log p}$.

Proof. Let $l_H(\boldsymbol{A}, r) = 2\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{A}\rangle - \frac{1}{a}\|\boldsymbol{X}\boldsymbol{A}\|_F^2 - aA_0[\rho^2 P_{2,H}(\boldsymbol{A};\frac{\lambda^o}{\rho}) + \sigma^2 r(m-r)]$, $l_0(\boldsymbol{A}, r)$ is defined similarly with $P_{2,H}$ replaced by $P_{2,0}$, $\mathcal{E}_H = \{\sup_{\boldsymbol{A}\in\boldsymbol{\Gamma}_{J,r}} l_H(\boldsymbol{A}, r) \geq at\sigma^2\}$, and $\mathcal{E}_0 = \{\sup_{\boldsymbol{A}:r(\boldsymbol{A})\leq r} l_0(\boldsymbol{A}, r) \geq at\sigma^2\}$.

Since $\mathcal{E}_H \subset \{\sup_{\boldsymbol{A}:r(A)\leq r} l_H(\boldsymbol{A}, r) \geq at\sigma^2\}$, the occurrence of $\mathcal{E}_H$ implies that $l_H(\boldsymbol{A}^o, r) \geq at\sigma^2$ for any

$$\boldsymbol{A}^o \in \arg\min_{r(\boldsymbol{A})\leq r} \frac{1}{a}\|\boldsymbol{X}\boldsymbol{A}\|_F^2 - 2\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{A}\rangle + aA_0\rho^2 P_{2,H}(\boldsymbol{A};\frac{\lambda^o}{\rho}).$$

From Lemma 2 and $a^2 A_0 \geq 2$, this further implies that there exists an $\boldsymbol{A}^o$ such that $l_0(\boldsymbol{A}^o, r) \geq at\sigma^2$. Hence $\mathcal{E}_H \subset \mathcal{E}_0$ and it suffices to show $\mathbb{P}(\mathcal{E}_0) \leq C\exp(-ct)$. Note that this reduction is by no means trivial because $P_0(s;\lambda) \geq P_H(s;\lambda), \forall s \in \mathbb{R}$.

Let $RS(\boldsymbol{A})$ be the row space of $\boldsymbol{A}$. Because

$$\begin{aligned} 2\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{A}\rangle &= 2\langle \mathbf{P}_{\boldsymbol{X}_{\mathcal{J}(\boldsymbol{A})}}\boldsymbol{E}, \boldsymbol{X}\boldsymbol{A}\rangle = 2\langle \mathbf{P}_{\boldsymbol{X}_{\mathcal{J}(\boldsymbol{A})}}\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{A})}, \boldsymbol{X}\boldsymbol{A}\rangle \\ &\leq 2\|\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}(\boldsymbol{A})}}\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{A})}\|_F\|\boldsymbol{X}\boldsymbol{A}\|_F \\ &\leq a\|\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}(\boldsymbol{A})}}\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{A})}\|_F^2 + \frac{1}{a}\|\boldsymbol{X}\boldsymbol{A}\|_F^2, \end{aligned}$$

we have $\mathbb{P}(\mathcal{E}_0) \leq \mathbb{P}(\sup_{\boldsymbol{A}:r(\boldsymbol{A})\leq r}\{\|\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}(\boldsymbol{A})}}\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{A})}\|_F^2 - A_0[\frac{1}{2}\sigma^2(r+\log p)J(\boldsymbol{A}) + \sigma^2 r(m-r)]\} \geq t\sigma^2)$. The conclusion follows from Lemma 3. □

Lemma 2. *Given any $K \geq \|\boldsymbol{X}\|_2^2$, there exists a globally optimal solution $\boldsymbol{B}^o$ to*

$$\min_{\boldsymbol{B}:\mathit{rank}(\boldsymbol{B})\leq r} \frac{1}{2K}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B}\|_F^2 + P_{2,H}(\boldsymbol{B};\lambda) \tag{23}$$

such that for any $j : 1 \leq j \leq p$, either $\boldsymbol{b}_j^o = \mathbf{0}$ or $\|\boldsymbol{b}_j^o\|_2 \geq \lambda$.

Proof. This follows similar lines as the proof of Theorem 2. Write $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{V}^T$ and define $F(\boldsymbol{S},\boldsymbol{V}) = \frac{1}{2K}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T\|_F^2 + P_{2,H}(\boldsymbol{S};\lambda)$. Let $\boldsymbol{B}_o = \boldsymbol{S}_o\boldsymbol{V}_o^T$ be any global minimizer. Then similar to Lemma 4, we can use a surrogate function $G(\boldsymbol{S},\tilde{\boldsymbol{S}}) = \|\boldsymbol{Y}\boldsymbol{V} - \boldsymbol{X}\tilde{\boldsymbol{S}}\|_F^2/(2K) + P_{2,H}(\tilde{\boldsymbol{S}};\lambda) + \langle(\boldsymbol{I} - \boldsymbol{X}^T\boldsymbol{X}/K)(\boldsymbol{S} - \tilde{\boldsymbol{S}}), \boldsymbol{S} - \tilde{\boldsymbol{S}}\rangle$ to prove that $\boldsymbol{S}^o = \vec{\Theta}_H(\frac{1}{K}\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}_o + (\boldsymbol{I} - \frac{1}{K}\boldsymbol{X}^T\boldsymbol{X})\boldsymbol{S}_o;\lambda)$ satisfies $F(\boldsymbol{S}_o,\boldsymbol{V}_o) \geq F(\boldsymbol{S}^o,\boldsymbol{V}_o)$, where $\vec{\Theta}_H$ is the multivariate version of the hard thresholding $\Theta_H(t;\lambda) = t1_{|t|>\lambda}$ (cf. Section 2.2). Note that to get this inequality we do not need the continuity assumption in Section 6.2. This means $\boldsymbol{B}^o = \boldsymbol{S}^o\boldsymbol{V}_o^T$ is also a global minimizer and $\|\boldsymbol{b}_j^o\|_2 = \|\boldsymbol{s}_j^o\|_2$ satisfy the desired property. □

Lemma 3. *Suppose all entries of $\boldsymbol{E}$ are iid $\sim \mathcal{N}(0,\sigma^2)$. For any given $1 \leq J \leq p, 1 \leq r \leq J$, define $R_{\mathcal{J},\boldsymbol{V}} := \|\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}}}\boldsymbol{E}\boldsymbol{V}\|_F^2 = \|\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}}}\boldsymbol{E}\mathbf{P}_{\boldsymbol{V}}\|_F^2$ where $|\mathcal{J}| = J$ and $\boldsymbol{V} \in \mathbb{O}^{m\times r}$, and $R_{J,r} := \sup_{|\mathcal{J}|=J,\boldsymbol{V}\in\mathbb{O}^{m\times r}} R_{\mathcal{J},\boldsymbol{V}}$. Let $P_o(J,r) = \sigma^2\{(q \wedge J + m - r)r + J\log(ep/J)\}$. Then the following result holds*

$$\mathbb{P}(R_{J,r} > t\sigma^2 + LP_o(J,r)) \leq C\exp(-ct), \forall t \geq 0 \tag{24}$$

where $L, C, c > 0$ are universal constants.

Proof. First, notice that for fixed $(\mathcal{J}, \boldsymbol{V})$, $R_{\mathcal{J},\boldsymbol{V}}/\sigma^2 \sim \chi^2(D)$ with $D \le (J \wedge q)r$. The standard tail bound for $X \sim \chi^2(D)$ gives $\mathbb{P}(X - D \ge t) \le \exp(-t^2/(4(D+t)))$, $\forall t \ge 0$ (Laurent and Massart, 2000a, Lemma 1).

Under the operator norm $\|\cdot\|_2$, let $\mathcal{T}$ be an ε-net of $\{\mathbf{P}_{\boldsymbol{V}} : \boldsymbol{V} \in \mathbb{O}^{m\times r}\}$ (which gives a Grassmann manifold). From Szarek (1982), its cardinality $\mathcal{N}(\varepsilon)$ satisfies $\mathcal{N}(\varepsilon) \le (C/\varepsilon)^{r(m-r)}$, where C is a universal constant. Now, for any $\boldsymbol{V}_1 \in \mathbb{O}^{m\times r}$, we can find a $\boldsymbol{V}_2 \in \mathbb{O}^{m\times r}$ with $\mathbf{P}_{\boldsymbol{V}_2} \in \mathcal{T}$ such that $\|\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T\|_2 \le \varepsilon$. Apply SVD: $[\boldsymbol{V}_1, \boldsymbol{V}_2] = [\tilde{\boldsymbol{U}}_1, \tilde{\boldsymbol{U}}_2]\tilde{\boldsymbol{D}}\tilde{V}^T$, where both $\tilde{\boldsymbol{U}}_1$ and $\tilde{\boldsymbol{U}}_2$ have r columns. Because $\|\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\boldsymbol{V}_1\|_F^2 = Tr\{\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\boldsymbol{V}_2\boldsymbol{V}_2^T\boldsymbol{E}^T\} + Tr\{\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}(\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T)\boldsymbol{E}^T\}$ and

$$\begin{aligned} Tr\{\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}(\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T)\boldsymbol{E}^T\} &= Tr\{\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\mathbf{P}_{\tilde{\boldsymbol{U}}}(\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T)\mathbf{P}_{\tilde{\boldsymbol{U}}}\boldsymbol{E}^T\} \\ &\le \varepsilon\|\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\mathbf{P}_{\tilde{\boldsymbol{U}}}\|_F^2 \\ &= \varepsilon\|\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\mathbf{P}_{\tilde{\boldsymbol{U}}_1}\|_F^2 + \varepsilon\|\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\mathbf{P}_{\tilde{\boldsymbol{U}}_2}\|_F^2, \end{aligned}$$

we obtain

$$R_{J,r} \le \max_{\mathcal{J}:|\mathcal{J}|=J, \mathbf{P}_{\boldsymbol{V}}\in\mathcal{T}} \|\mathbf{P}_{\boldsymbol{X}_\mathcal{J}}\boldsymbol{E}\mathbf{P}_{\boldsymbol{V}}\|_F^2 + 2\varepsilon R_{J,r}.$$

Applying the χ^2 tail bound and the union bound gives

$$\mathbb{P}\left(R_{J,r} \ge \sigma^2\frac{(J\wedge q)r + t}{1-2\varepsilon}\right) \le \binom{p}{J}\left(\frac{C}{\varepsilon}\right)^{r(m-r)}\exp(-Ct^2/(t + (J\wedge q)r).$$

Hence we obtain for some L large enough, (24) holds for any $t \ge 0$. □

In summary, the stochastic term can be bounded by

$$\begin{aligned} 2\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle \le& \frac{1}{a}\|\boldsymbol{X}\boldsymbol{\Delta}\|_F^2 + aA_0[\rho^2 P_{2,H}(\boldsymbol{\Delta}; \frac{\lambda^o}{\rho}) + 2\sigma^2 r(m-r)] + R \\ \le& \frac{2}{a}\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + \frac{2}{a}\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \\ &+ 2aA_0\sigma^2 r(m-r) + R + aA_0\rho^2 P_{2,H}(\boldsymbol{\Delta}; \frac{\lambda^o}{\rho}), \end{aligned}$$

where $R := \sup_{1\le J\le p, 1\le r\le m\wedge J}\sup_{\boldsymbol{A}\in\boldsymbol{\Gamma}_{J,r}} l(\boldsymbol{A}, r)$. From Lemma 1, it is not difficult to show that $\mathbb{E}R \le aC\sigma^2$.

Substituting the bound in (21) gives

$$\begin{aligned} (1-\frac{2}{a})\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \le (1+\frac{2}{a})\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + 2aA_0\sigma^2 r(m-r) + R \\ +2\rho^2(P(\boldsymbol{B}; \frac{\lambda}{\rho}) - P(\hat{\boldsymbol{B}}; \frac{\lambda}{\rho})) + aA_0\rho^2 P_{2,H}(\boldsymbol{\Delta}; \frac{\lambda^o}{\rho}). \end{aligned}$$

It remains to deal with $2\rho^2(P(\boldsymbol{B};\frac{\lambda}{\rho}) - P(\hat{\boldsymbol{B}};\frac{\lambda}{\rho})) + aA_0\rho^2 P_{2,H}(\boldsymbol{\Delta};\frac{\lambda^o}{\rho}) =: I$.

When $P = P_0$, $\rho^2 P_{2,H}(\boldsymbol{\Delta};\frac{\lambda^o}{\rho}) \le \rho^2 P_{2,0}(\boldsymbol{\Delta};\frac{\lambda^o}{\rho}) \le \rho^2 P_{2,0}(\boldsymbol{B};\frac{\lambda^o}{\rho}) + \rho^2 P_{2,0}(\hat{\boldsymbol{B}};\frac{\lambda^o}{\rho})$. The oracle bound thus follows, by choosing $a > 2 \vee \sqrt{2/A_0}$ and $\lambda = A\lambda^o$ with $A > \sqrt{aA_0/2}$.

When $P = P_1$, $I \le 2\rho^2(P_{2,1}(\boldsymbol{B};\frac{\lambda}{\rho}) - P_{2,1}(\hat{\boldsymbol{B}};\frac{\lambda}{\rho})) + 2aA_0\rho^2 P_{2,1}(\boldsymbol{\Delta};\frac{\lambda^o}{\rho}) = \rho\lambda^o(2A\|\boldsymbol{B}\|_{2,1} - 2A\|\hat{\boldsymbol{B}}\|_{2,1} + aA_0\|\boldsymbol{\Delta}\|_{2,1})$ under $\lambda = A\lambda^o$. Let $\mathcal{J} = \mathcal{J}(\boldsymbol{B})$ and $J = J(\boldsymbol{B})$. Using the sub-additivity of P_1, we have $I \le \rho\lambda^o\{(2A + aA_0)\|\boldsymbol{\Delta}_{\mathcal{J}}\|_{2,1} - (2A - aA_0)\|\boldsymbol{\Delta}_{\mathcal{J}^c}\|_{2,1}\}$.

Apply a comparison trick. The case $(1+\frac{2}{a})\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + aA_0\sigma^2 r(m-r) + R \ge \rho\lambda^o\{(2A+aA_0)\|\boldsymbol{\Delta}_{\mathcal{J}}\|_{2,1} - (2A-aA_0)\|\boldsymbol{\Delta}_{\mathcal{J}^c}\|_{2,1}\}$ is trivial. Suppose the reverse inequality holds, which implies that $\boldsymbol{\Delta}$ falls into a cone $\{\boldsymbol{M} : \|\boldsymbol{M}_{\mathcal{J}^c}\|_{2,1} \le (1 + 2aA_0/(2A - aA_0))\|\boldsymbol{M}_{\mathcal{J}}\|_{2,1}\}$. Moreover, by Cauchy-Schwarz inequality, $2\rho\lambda^o\|\boldsymbol{\Delta}_{\mathcal{J}}\|_{2,1} \le \lambda^{o2} Jb + \rho^2\|\boldsymbol{\Delta}_{\mathcal{J}}\|_F^2/b$ for any $b > 0$.

We introduce a multivariate restrictive eigenvalue assumption:

ASSUMPTION $\mathfrak{A}$. We say $\boldsymbol{\Sigma} \in \mathbb{R}^{p\times p}$ satisfies condition $\mathfrak{A}(\mathcal{J}, \delta_{\mathcal{J}}, \vartheta)$ for an index set $\mathcal{J} \subseteq [p]$ and positive number $\delta_{\mathcal{J}}$, if and only if

$$Tr(\boldsymbol{M}^T\boldsymbol{\Sigma}\boldsymbol{M}) \ge \delta_{\mathcal{J}} \sum_{j\in\mathcal{J}} \|\boldsymbol{m}_j\|_2^2 \tag{25}$$

for all $p \times m$ matrices $\boldsymbol{M}$ (with rows $\boldsymbol{m}_j$) satisfying $(1+\vartheta)\sum_{j\in\mathcal{J}}\|\boldsymbol{m}_j\|_2 \ge \sum_{j\in\mathcal{J}^c}\|\boldsymbol{m}_j\|_2$ where $\vartheta > 0$ is a constant.

Under the condition that $\boldsymbol{\Sigma} = \boldsymbol{X}^T\boldsymbol{X}/\rho^2$ satisfies Assumption $\mathfrak{A}(\mathcal{J}(\boldsymbol{B}), \delta_{\mathcal{J}(\boldsymbol{B})}, \vartheta)$ with $\delta_{\mathcal{J}(\boldsymbol{B})}, \vartheta$ being positive constants, we obtain

$$\begin{aligned}
&(1-\frac{2}{a})\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \\
\le&(2A + aA_0)(\lambda^{o2} Jb + \rho^2\|\boldsymbol{\Delta}_{\mathcal{J}}\|_F^2/b) \\
\le&(2A + aA_0)\lambda^{o2} Jb + (2A + aA_0)\|\boldsymbol{X}\boldsymbol{\Delta}\|_F^2/(b\delta_{\mathcal{J}}) \\
\le&(2A + aA_0)\lambda^{o2} Jb + 2(2A + aA_0)\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2/(b\delta_{\mathcal{J}}) \\
&+ 2(2A + aA_0)\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2/(b\delta_{\mathcal{J}}).
\end{aligned}$$

Then choosing the constants $a > 2\vee\sqrt{2/A_0}$, $A \ge aA_0(\frac{1}{2} + \frac{1}{\vartheta})$, and $b > \frac{2(2A+aA_0)}{(1-2/a)\delta_{\mathcal{J}}}$, we obtain the oracle inequality.

Remark. In the group ℓ_1 case, we did not bound $\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle$ in the more or less standard way: $\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle = \langle \boldsymbol{X}^T\boldsymbol{E}, \boldsymbol{\Delta}\rangle = \sum_{j=1}^p \boldsymbol{x}_j^T\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{\Delta})}(\hat{\boldsymbol{b}}_j - \boldsymbol{b}_j) \le \max_{1\le j\le p}\|\boldsymbol{x}_j^T\boldsymbol{E}\mathbf{P}_{RS(\boldsymbol{\Delta})}\|_2 \cdot \|\hat{\boldsymbol{B}} - \boldsymbol{B}\|_{2,1}$, where $\boldsymbol{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_p]$, which, for the

choice of $\lambda = A\sigma\sqrt{r+\log p}$, gives a suboptimal error bound (with high probability).

6.2 Proof of Theorem 2

Recall the objective function $F(\boldsymbol{B};\lambda)$ in (7), or $F(\boldsymbol{s}_1,\cdots,\boldsymbol{s}_p,\boldsymbol{V};\lambda) = F(\boldsymbol{S},\boldsymbol{V};\lambda) = \|\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{S}\boldsymbol{V}^T\|_F^2/(2K) + \sum_j P(\|\boldsymbol{s}_j\|_2);\lambda)$ with $\boldsymbol{S} = [\boldsymbol{s}_1,\cdots,\boldsymbol{s}_p]^T$. Given $\boldsymbol{V}$, the optimization over $\boldsymbol{S}$ reduces to a group variable selection problem

$$\frac{1}{2K}\|\operatorname{vec}((\boldsymbol{Y}\boldsymbol{V})^T) - (\boldsymbol{X}\otimes\boldsymbol{I})\operatorname{vec}(\boldsymbol{S}^T)\|_F^2 + \sum P(\|\boldsymbol{s}_j\|_2;\lambda), \qquad (26)$$

where $\otimes$ is the Kronecker product. Define $T_{\boldsymbol{V}}:\mathbb{R}^{p\times r}\to\mathbb{R}^{p\times r}$ by $T_{\boldsymbol{V}}\circ\boldsymbol{S} = \vec{\Theta}(\boldsymbol{\xi};\lambda)$ with $\boldsymbol{\xi} = [\boldsymbol{\xi}_1,\cdots,\boldsymbol{\xi}_p]^T = \boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}/K + (\boldsymbol{I}-\boldsymbol{X}^T\boldsymbol{X}/K)\boldsymbol{S}$. Applying Theorem 1 in She (2012) we get

Lemma 4. *Let Θ be an arbitrary threshold function satisfying Definition 2.1 and P be an associated penalty given by (8). Given any $\boldsymbol{V}\in\mathbb{O}^{m\times r}$ and $\boldsymbol{S}\in\mathbb{R}^{p\times r}$, define $\tilde{\boldsymbol{S}} = T_{\boldsymbol{V}}\circ\boldsymbol{S} = \vec{\Theta}(\boldsymbol{\xi};\lambda)$. If $K\ge\|\boldsymbol{X}\|_2^2$ and Θ is continuous at $\|\boldsymbol{\xi}_1\|_2,\ldots,\|\boldsymbol{\xi}_p\|_2$, then $F(\boldsymbol{S},\boldsymbol{V}) - F(\tilde{\boldsymbol{S}},\boldsymbol{V}) \ge \frac{1-\|\boldsymbol{X}\|_2^2/K}{2}\|\tilde{\boldsymbol{S}}-\boldsymbol{S}\|_F^2$.*

Back to the algorithm, let $\boldsymbol{\xi}^{(t,t')} = \frac{1}{K}\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{V}^{(t-1)} + (\boldsymbol{I}-\frac{1}{K}\boldsymbol{X}^T\boldsymbol{X})\tilde{\boldsymbol{S}}^{(t'-1)}$. The following **continuity assumption** is made throughout the proof:

Assumption $\mathfrak{B}$: $\vec{\Theta}$ is continuous at any point in the closure of $\{\boldsymbol{\xi}^{(t,t')}\}$.

For continuous thresholding rules such as soft-thresholding, this regularity condition always holds. Practically used thresholding rules (such as hard-thresholding) have few discontinuity points and such discontinuities rarely occur in real applications.

For the inner loop in Algorithm 1, we have $F(\boldsymbol{S}^{(t-1)},\boldsymbol{V}^{(t)}) - F(\boldsymbol{S}^{(t)},\boldsymbol{V}^{(t)}) \ge \frac{(1-\|\boldsymbol{X}\|_2^2/K)}{2}\sum_{t'}\|\tilde{\boldsymbol{S}}^{(t'-1)}-\tilde{\boldsymbol{S}}^{(t')}\|_F^2 \ge \frac{(1-\|\boldsymbol{X}\|_2^2/K)}{2}\|\boldsymbol{S}^{(t-1)}-\boldsymbol{S}^{(t)}\|_F^2$, based on the triangle inequality. Given $\boldsymbol{S}$, the Procrustes rotation performed in Steps (3)-(4) gives a globally optimal solution over all $\boldsymbol{V}\in\mathbb{O}^{m\times r}$ (cf. She (2013)). Hence $F(\boldsymbol{S}^{(t-1)},\boldsymbol{V}^{(t-1)}) \ge F(\boldsymbol{S}^{(t-1)},\boldsymbol{V}^{(t)}) \ge F(\boldsymbol{S}^{(t)},\boldsymbol{V}^{(t)}) + \frac{(1-\|\boldsymbol{X}\|_2^2/K)}{2}\|\boldsymbol{S}^{(t-1)}-\boldsymbol{S}^{(t)}\|_F^2$, $\forall t\ge 1$.

To prove part (ii), we first observe that when updating $\boldsymbol{S}$, the algorithm allows one to perform $T_{\boldsymbol{V}^{(t)}}$ any times (denoted by α_t) provided α_t does not go beyond

$M_{inner} \in \mathbb{N}$ that is pre-specified before running the algorithm. To handle the ambiguity of $\boldsymbol{S}^{(t)}$ and the non-uniqueness in $\boldsymbol{V}$-optimization, we use the tool of *point-to-set* mappings (Luenberger and Ye, 2008). Let $T_{\boldsymbol{V}}^{\alpha}$ ($\alpha \in \mathbb{N}$) be the composition of α $T_{\boldsymbol{V}}$'s. Define point-to-set maps $\mathcal{M}^S(\boldsymbol{S}, \boldsymbol{V}) = \{(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}}) \in \Omega : \bar{\boldsymbol{S}} \in \{T_{\boldsymbol{V}} \circ \boldsymbol{S}, T_{\boldsymbol{V}}^2 \circ \boldsymbol{S}, \cdots, T_{\boldsymbol{V}}^{M_{inner}} \circ \boldsymbol{S}\}, \bar{\boldsymbol{V}} = \boldsymbol{V}\}$, and $\mathcal{M}^V(\boldsymbol{S}, \boldsymbol{V}) = \{(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}}) \in \Omega : \inf_{\tilde{\boldsymbol{V}} \in \mathbb{O}^{m \times r}} F(\boldsymbol{S}, \tilde{\boldsymbol{V}}) \geq F(\boldsymbol{S}, \bar{\boldsymbol{V}}), \bar{\boldsymbol{S}} = \boldsymbol{S}\}$. Then we can use $\mathcal{M} := \mathcal{M}^S \mathcal{M}^V$ to characterize Algorithm 1.

By the continuity assumption, it is easy to show that $\mathcal{M}$ is a *closed* at any accumulation point of $\{\boldsymbol{S}^{(t-1)}, \boldsymbol{V}^{(t)}\}$ (cf. Lemma 14 in Bunea et al. 2012). From part (i), we have $F(\boldsymbol{S}^{(0)}, \boldsymbol{V}^{(1)}) - F(\boldsymbol{S}^{(t)}, \boldsymbol{V}^{(t)}) \geq \frac{(1-\|\boldsymbol{X}\|_2^2/K)}{2} \|\boldsymbol{S}^{(0)} - \boldsymbol{S}^{(t)}\|_F^2$, and so $\|\boldsymbol{S}^{(t)}\|_2 \leq \sqrt{\frac{2F(\boldsymbol{S}^{(0)}, \boldsymbol{V}^{(1)})}{1-\|\boldsymbol{X}\|_2^2/K}} + \|\boldsymbol{S}^{(0)}\|_2$. That is, $\boldsymbol{S}^{(t)}$, $\boldsymbol{V}^{(t)}$, $\boldsymbol{B}^{(t)}$ are all uniformly bounded in t.

The rest of the proof proceeds along similar lines of the proof of Theorem 7 in Bunea et al. (2012) with minor modifications. We give some details as follows. Suppose $(\boldsymbol{S}^*, \boldsymbol{V}^*)$ is an arbitrary accumulation point of $(\boldsymbol{S}^{(t-1)}, \boldsymbol{V}^{(t)})$ with $(\boldsymbol{S}^*, \boldsymbol{V}^*) \in \Omega$. From the boundedness of the sequence, assume without loss of generality that $(\boldsymbol{S}^{(t)}, \boldsymbol{V}^{(t+1)})$ converges to $(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}})$. Then $(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}}) \in \mathcal{M}(\boldsymbol{S}^*, \boldsymbol{V}^*)$ because $\mathcal{M}$ is closed. It follows from part (i) that $F(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}}) = F(\boldsymbol{S}^*, \boldsymbol{V}^*)$, and thus $F(\bar{\boldsymbol{S}}, \boldsymbol{V}^*) = F(\boldsymbol{S}^*, \boldsymbol{V}^*)$. Lemma 4 implies that $\bar{\boldsymbol{S}} = \boldsymbol{S}^*$, that is, $T_{\boldsymbol{V}^*}^{\alpha_0} \circ \boldsymbol{S}^* = \boldsymbol{S}^*$ for some $\alpha_0 \in \mathbb{N}$. We have $F(T_{\boldsymbol{V}^*}^{\alpha} \circ \boldsymbol{S}^*, \boldsymbol{V}^*) = F(\boldsymbol{S}^*, \boldsymbol{V}^*)$ for any $\alpha \leq \alpha_0$, and in particular, $F(T_{\boldsymbol{V}^*} \circ \boldsymbol{S}^*, \boldsymbol{V}^*) = F(\boldsymbol{S}^*, \boldsymbol{V}^*)$. Applying Lemma 4 again yields $T_{\boldsymbol{V}^*} \circ \boldsymbol{S}^* = \boldsymbol{S}^*$, i.e., $\boldsymbol{S}^*$ is a fixed point of $T_{\boldsymbol{V}^*}$. Following the lines of the proof of Lemma 1 in She (2012), it can be shown that for fixed $(\boldsymbol{s}_1, \cdots, \boldsymbol{s}_{j-1}, \boldsymbol{s}_{j+1}, \cdots, \boldsymbol{s}_p, \boldsymbol{V}^*)$, $\boldsymbol{s}_j^*$ is a global minimizer of $F(\boldsymbol{s}_1, \cdots, \boldsymbol{s}_j, \cdots, \boldsymbol{s}_p, \boldsymbol{V}^*)$ under the continuity assumption, $1 \leq j \leq p$. On the other hand, from $\bar{\boldsymbol{S}} = \boldsymbol{S}^*$, we have $(\boldsymbol{S}^*, \bar{\boldsymbol{V}}) \in \mathcal{M}^V(\boldsymbol{S}^*, \boldsymbol{V}^*)$. $\bar{\boldsymbol{V}}$ is a (global) minimizer of $F(\boldsymbol{S}^*, \boldsymbol{V})$ given $\boldsymbol{S}^*$. But $F(\boldsymbol{S}^*, \bar{\boldsymbol{V}}) = F(\bar{\boldsymbol{S}}, \bar{\boldsymbol{V}}) = F(\boldsymbol{S}^*, \boldsymbol{V}^*)$, so $\boldsymbol{V}^*$ also minimizes $F(\boldsymbol{S}^*, \boldsymbol{V})$ given $\boldsymbol{S}^*$. The proof is complete.

6.3 Proof of Theorem 3

Given any matrix $\boldsymbol{A}$, we use $CS(\boldsymbol{A})$ and $RS(\boldsymbol{A})$ to denote its column space and row space, respectively. Denote by $\mathbf{P}_{\boldsymbol{A}}$ the orthogonal projection matrix onto $CS(\boldsymbol{A})$, i.e., $\mathbf{P}_{\boldsymbol{A}} = \boldsymbol{A}(\boldsymbol{A}^T\boldsymbol{A})^+\boldsymbol{A}^T$, where $^+$ stands for the Moore-Penrose pseudoinverse, and $\mathbf{P}_{\boldsymbol{A}}^{\perp}$ the projection onto its orthogonal complement. Given any index set $\mathcal{J} \subset [p]$, $\mathbf{P}_{\mathcal{J}}$ is short for $\mathbf{P}_{\boldsymbol{X}_{\mathcal{J}}}$ when there is no ambiguity.

In the theorem, the noise matrix has sub-Gaussian marginal tails.

Definition 6.1. *ξ is called a sub-Gaussian random variable if there exist constants $C, c > 0$ such that $\mathbb{P}\{|\xi| \geq t\} \leq Ce^{-ct^2}, \forall t > 0$. The scale ($\psi_2$-norm) for ξ is defined as $\sigma(\xi) = \inf\{\sigma > 0 : \mathbb{E}\exp(\xi^2/\sigma^2) \leq 2\}$. $\boldsymbol{\xi} \in \mathbb{R}^p$ is called a sub-Gaussian random vector with scale bounded by σ if all one-dimensional marginals $\langle \boldsymbol{\xi}, \boldsymbol{\alpha}\rangle$ are sub-Gaussian satisfying $\|\langle \boldsymbol{\xi}, \boldsymbol{\alpha}\rangle\|_{\psi_2} \leq \sigma\|\boldsymbol{\alpha}\|_2, \forall \boldsymbol{\alpha} \in R^p$.*

Sub-Gaussian examples include Gaussian random variables and bounded random variables such as Bernoulli. In the proof, we assume vec $(\boldsymbol{E})$ is sub-Gaussian. Note that the entries of $\boldsymbol{E}$ may not be iid.

Recall $P_o(\boldsymbol{B}) = \sigma^2\{(q \wedge J(\boldsymbol{B}) + m - r(\boldsymbol{B}))r(\boldsymbol{B}) + J(\boldsymbol{B})\log(ep/J(\boldsymbol{B}))\}$ with $q = \text{rank}(\boldsymbol{X})$; for convenience, we also denote it by $P_o(J(\boldsymbol{B}), r(\boldsymbol{B}))$. From the construction of $\hat{\boldsymbol{B}}$, we have

$$\frac{1}{2}\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \leq \frac{1}{2}\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + AP_o(\boldsymbol{B}) - AP_o(\hat{\boldsymbol{B}}) + \langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}\rangle \tag{27}$$

Let $\boldsymbol{\Delta} = \hat{\boldsymbol{B}} - \boldsymbol{B}$, $\hat{\mathcal{J}} = \mathcal{J}(\hat{\boldsymbol{B}})$, $\mathcal{J} = \mathcal{J}(\boldsymbol{B})$, $J = J(\boldsymbol{B})$, $\hat{J} = J(\hat{\boldsymbol{B}})$, $r = r(\boldsymbol{B})$, $\hat{r} = r(\hat{\boldsymbol{B}})$. Let $\mathbf{P}_{rs}$ and $\mathbf{P}_{\mathcal{J}}$ be the orthogonal projections onto the row space of $\boldsymbol{X}_{\mathcal{J}}\boldsymbol{B}_{\mathcal{J}}$ and the column space of $\boldsymbol{X}_{\mathcal{J}}$, respectively. Decompose $\boldsymbol{X}\boldsymbol{\Delta}$ as follows

$$\begin{aligned}\boldsymbol{X}\boldsymbol{\Delta} =& \boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs} + \boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}^{\perp}\\ =& \mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs} + \mathbf{P}_{\mathcal{J}}^{\perp}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs} + \boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}^{\perp}\\ =& \mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs} + \mathbf{P}_{\mathcal{J}}^{\perp}\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs} + \boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}^{\perp}.\end{aligned}$$

Clearly, $\|\boldsymbol{X}\boldsymbol{\Delta}\|_F^2 = \|\mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}\|_F^2 + \|\mathbf{P}_{\mathcal{J}}^{\perp}\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}\|_F^2 + \|\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}^{\perp}\|_F^2$. The noise term in (27) is now

$$\begin{aligned}\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle =& \langle \boldsymbol{E}, \mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}\rangle + \langle \boldsymbol{E}, \mathbf{P}_{\mathcal{J}}^{\perp}\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}\rangle + \langle \boldsymbol{E}, \boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}^{\perp}\rangle\\ \equiv& I + II + III.\end{aligned} \tag{28}$$

To bound *I* and *III*, we introduce Lemma 5.

Lemma 5. *Suppose vec $(\boldsymbol{E})$ is sub-Gaussian with mean zero and ψ_2-norm bounded by σ. Given $\boldsymbol{X} \in \mathbb{R}^{n\times p}$, $1 \leq J \leq p$, $1 \leq r \leq J \wedge m$, define $\Gamma_{J,r} = \{\boldsymbol{A} \in \mathbb{R}^{n\times m} : \|\boldsymbol{A}\|_F \leq 1, rank(\boldsymbol{A}) \leq r, CS(\boldsymbol{A}) \subset CS(\boldsymbol{X}_{\mathcal{J}})$ for some $\mathcal{J} : |\mathcal{J}| = J\}$. Let $P_o'(J, r) = \sigma^2\{(q \wedge J)r + (m - r)r + \log\binom{p}{J}\}$. Then for any $t \geq 0$,*

$$\mathbb{P}\left(\sup_{\boldsymbol{A}\in\Gamma_{J,r}} \langle \boldsymbol{E}, \boldsymbol{A}\rangle \geq t\sigma + \sqrt{L \cdot P_o'(J, r)}\right) \leq C\exp(-ct^2), \tag{29}$$

where $L, C, c > 0$ are universal constants.

Proof. By Definition 6.1, for any fixed $\boldsymbol{A}$, $\langle \boldsymbol{E}, \boldsymbol{A}\rangle$ is a mean-centered sub-Gaussian random variable with scale bounded by $\sigma\|\boldsymbol{A}\|_F$. Therefore, $\{\langle \boldsymbol{E}, \boldsymbol{A}\rangle : \boldsymbol{A} \in \Gamma_{J,r}\}$ is a stochastic process with sub-Gaussian increments; Dudley's entropy integral can be used to bound its supremum, see, e.g., Talagrand (2005). The induced metric on $\Gamma_{J,r}$ is Euclidean: $d(\boldsymbol{A}_1, \boldsymbol{A}_2) = \sigma\|\boldsymbol{A}_1 - \boldsymbol{A}_2\|_F$.

To compute the metric entropy $\log \mathcal{N}(\varepsilon, \Gamma_{J,r}, d)$, where $\mathcal{N}(\varepsilon, \Gamma_{J,r}, d)$ is the smallest cardinality of an ε-net that covers $\Gamma_{J,r}$ under d, we characterize each matrix in $\Gamma_{J,r}$ using its row/column spaces, motivated by Recht et al. (2010). Given $\boldsymbol{A} \in \Gamma_{J,r}$, its column space must be contained in $CS(\boldsymbol{X}_{\mathcal{J}})$ for some $\mathcal{J}$ with $|\mathcal{J}| = J$; its row space must be contained in an r-dimensional subspace in $\mathbb{R}^m$. Hence

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T, \tag{30}$$

where $\boldsymbol{\Sigma} \in \mathbb{R}^{(J\wedge q)\times r}$ and $\mathbf{P}_{\boldsymbol{U}} = \mathbf{P}_{\boldsymbol{X}_{\mathcal{J}}}$. $\boldsymbol{\Sigma}$ is in a $(J \wedge q) \times r$-dimensional unit ball (denoted by $B_{(J\wedge q)\times r}$).

The number of candidate $\boldsymbol{X}_{\mathcal{J}}$ is $\binom{p}{J}$. By a standard volume argument,

$$\mathcal{N}(\varepsilon, B_{(J\wedge q)\times r}, d') \le (C_0/\varepsilon)^{(J\wedge q)\times r},$$

where d' is the Euclidean distance in $\mathbb{R}^{(J\wedge q)\times r}$ and C_0 is a universal constant. Note that $RS(\boldsymbol{A})$ is a point on the Grassmann manifold (denoted by $G_{m,r}$) of all r-dimensional subspaces of $\mathbb{R}^m$. Equipped with the metric d'' given by the operator norm $\|\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T\|_2$ for any $\boldsymbol{V}_1, \boldsymbol{V}_2 \in \mathbb{O}^{m\times r}$, we have $\mathcal{N}(\varepsilon, G_{m,r}, d'') \le (C_1/\varepsilon)^{r(m-r)}$, where C_1 is a universal constant (Szarek, 1982). We claim that

$$\mathcal{N}(\varepsilon, \Gamma_{J,r}, d) \le \binom{p}{J}\left(\frac{C\sigma}{\varepsilon}\right)^{(J\wedge q)\times r + r(m-r)}. \tag{31}$$

In fact, given any $\boldsymbol{A}_1 \in \Gamma_{J,r}$, with $\boldsymbol{A}_1 = \boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T$ according to (30), find $\boldsymbol{V}_2$ and $\boldsymbol{\Sigma}_2$ such that $\|\boldsymbol{V}_1\boldsymbol{V}_1^T - \boldsymbol{V}_2\boldsymbol{V}_2^T\|_2 \le \varepsilon$ and $\|\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T\boldsymbol{V}_2 - \boldsymbol{\Sigma}_2\|_F \le \varepsilon$, then, for $\boldsymbol{A}_2 = \boldsymbol{U}_1\boldsymbol{\Sigma}_2\boldsymbol{V}_2^T$,

$$\begin{aligned}
\|\boldsymbol{A}_1 - \boldsymbol{A}_2\|_F &\le \|\boldsymbol{A}_1 - \boldsymbol{A}_1\boldsymbol{V}_2\boldsymbol{V}_2^T\|_F + \|\boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T\boldsymbol{V}_2\boldsymbol{V}_2^T - \boldsymbol{U}_1\boldsymbol{\Sigma}_2\boldsymbol{V}_2^T\|_F \\
&\le (Tr\{\boldsymbol{A}_1^T\boldsymbol{A}_1(\mathbf{P}_{\boldsymbol{V}_1} - \mathbf{P}_{\boldsymbol{V}_2})^2\})^{1/2} + \|\boldsymbol{\Sigma}_1\boldsymbol{V}_1^T\boldsymbol{V}_2 - \boldsymbol{\Sigma}_2\|_F \\
&\le (\|\boldsymbol{A}_1\|_F^2\|\mathbf{P}_{\boldsymbol{V}_1} - \mathbf{P}_{\boldsymbol{V}_2}\|_2^2)^{1/2} + \varepsilon \le 2\varepsilon.
\end{aligned}$$

Using Dudley's integral bound, we obtain

$$\mathbb{P}\left(\sup_{\boldsymbol{A}\in\Gamma_{J,r}} \langle \boldsymbol{E}, \boldsymbol{A}\rangle \ge t\sigma + L\int_0^\sigma \sqrt{\log \mathcal{N}(\varepsilon, \Gamma_{J,r}, d)}\,\mathrm{d}\varepsilon\right) \le C\exp(-ct^2).$$

Simple computation yields

$$\int_0^\sigma \sqrt{\log \mathcal{N}(\varepsilon, \Gamma_{J,r}, d)}\, \mathrm{d}\varepsilon) \lesssim \sigma\sqrt{\log \binom{p}{J}} + \sigma\sqrt{(J \wedge q) \times r + r(m-r)} \lesssim \sqrt{P_o'(J,r)}$$

due to the Cauchy-Schwarz inequality. □

Terms *I* and *III* on the right hand side of (28) can handled in a similar manner. For instance,

$$\begin{aligned}
&\langle \boldsymbol{E}, \mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs} \rangle - \frac{1}{a} \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F^2 - bLP_o(J,r) \\
\leq & \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F \langle \boldsymbol{E}, \mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs} / \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F \rangle - 2\sqrt{\frac{b}{a}} \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F \sqrt{LP_o(J,r)} \\
\leq & \frac{1}{a'} \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F^2 + \frac{a'}{4} \sup_{1 \leq J \leq p, 1 \leq r \leq m \wedge J} \left(\sup_{\boldsymbol{A} \in \Gamma_{J,r}} \langle \boldsymbol{E}, \boldsymbol{A} \rangle - 2\sqrt{\frac{b}{a}} \sqrt{LP_o(J,r)} \right)_+^2 \\
\equiv & \frac{1}{a'} \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F^2 + \frac{a'}{4} \sup_{1 \leq J \leq p, 1 \leq r \leq m \wedge J} R_{J,r}^2 \\
\equiv & \frac{1}{a'} \|\mathbf{P}_{\mathcal{J}} \boldsymbol{X} \boldsymbol{\Delta} \mathbf{P}_{rs}\|_F^2 + \frac{a'}{4} R^2,
\end{aligned}$$

for any $a, b, a' > 0$. Lemma 5 provides a control of R^2:

$$\begin{aligned}
& \mathbb{P}(R \geq t\sigma) \\
\leq & \sum_{J=1}^{p} \sum_{r=1}^{m \wedge J} \mathbb{P}(R_{J,r} \geq t\sigma) \\
\leq & \sum_{J=1}^{p} \sum_{r=1}^{m \wedge J} \mathbb{P}(\sup_{\boldsymbol{A} \in \Gamma_{J,r}} \langle \boldsymbol{E}, \boldsymbol{A} \rangle - \sqrt{LP_o(J,r)} \geq t\sigma + (2\sqrt{b/a} - 1)\sqrt{LP_o(J,r)}) \\
\leq & \sum_{J=1}^{p} \sum_{r=1}^{m \wedge J} C \exp(-ct^2) \exp\left\{ -c \left((2\sqrt{b/a} - 1)^2 L \cdot P_o(J,r)/\sigma^2 \right) \right\} \\
\leq & C' \exp(-ct^2),
\end{aligned}$$

from which it follows that $\mathbb{E}R^2 \leq C\sigma^2$.

Similarly, we introduce the following lemma to deal with *II*.

Lemma 6. *Suppose vec* $(\boldsymbol{E})$ *is sub-Gaussian with mean zero and* ψ_2*-norm bounded by* σ*. Given* $\boldsymbol{X} \in \mathbb{R}^{n\times p}$*,* $1 \le J, J' \le p$*,* $1 \le r \le J \wedge m$*, define* $\Gamma_{J',J,r} = \{\boldsymbol{A} \in \mathbb{R}^{n\times m} : \|\boldsymbol{A}\|_F \le 1, rank(\boldsymbol{A}) \le r, CS(\boldsymbol{A}) \subset CS(\mathbf{P}^{\perp}_{\mathcal{J}'}\mathbf{P}_{\mathcal{J}})$ *for some* $\mathcal{J}', \mathcal{J} \subset [p]$ *satisfying* $|\mathcal{J}'| = J', |\mathcal{J}| = J\}$*. Let* $P''_o(J', J, r) = \sigma^2\{(q \wedge J \wedge (p - J'))r + (m - r)r + \log\binom{p}{J} + \log\binom{p}{J'}\}$*. Then for any* $t \ge 0$*,*

$$\mathbb{P}\left(\sup_{\boldsymbol{A}\in\Gamma_{J',J,r}} \langle \boldsymbol{E}, \boldsymbol{A}\rangle \ge t\sigma + \sqrt{L \cdot P''_o(J', J, r)}\right) \le C\exp(-ct^2), \tag{32}$$

where $L, C, c > 0$ *are universal constants.*

Proof. The proof is similar to that of Lemma 5 based on the entropy integral bound. The rate P''_o can be obtained by calculating the degrees-of-freedom and inflation bounds. The details are omitted. □

Similar to the treatment of I, applying Lemma 6 gives (noticing that $P''_o(J, \hat{J}, \hat{r}) \le P_o(J, r) + P_0(\hat{J}, \hat{r})$):

$$\begin{aligned}
&\langle \boldsymbol{E}, \mathbf{P}^{\perp}_{\hat{\mathcal{J}}}\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}\rangle - \frac{1}{a}\|\mathbf{P}^{\perp}_{\hat{\mathcal{J}}}\boldsymbol{X}_{\hat{\mathcal{J}}}\hat{\boldsymbol{B}}_{\hat{\mathcal{J}}}\mathbf{P}_{rs}\|_F^2 - bL\{P_o(J, r) + P_0(\hat{J}, \hat{r})\}\\
\le&\frac{1}{a'}\|\mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}\|_F^2 + \frac{a'}{4}R'^2,
\end{aligned}$$

with $\mathbb{E}R'^2 \le C\sigma^2$.

In summary, for any $a, b, a', b' > 0$, we have

$$\begin{aligned}
&\langle \boldsymbol{E}, \boldsymbol{X}\boldsymbol{\Delta}\rangle\\
\le&(\frac{1}{a} + \frac{1}{a'})\|\boldsymbol{X}\boldsymbol{\Delta}\|_F^2 + bL\{P_o(J, r) + P_o(J, r) + P_0(\hat{J}, \hat{r}) + P_o(\hat{J}, \hat{r})\}\\
&+ \frac{a'}{4}(2R^2 + R'^2)\\
\le&(\frac{1}{a} + \frac{1}{a'})(1 + b')\|\boldsymbol{X}\boldsymbol{B} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + (\frac{1}{a} + \frac{1}{a'})(1 + \frac{1}{b'})\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2\\
&+ 2bL\{P_o(J, r) + P_o(\hat{J}, \hat{r})\} + \frac{a'}{4}(2R^2 + R'^2).
\end{aligned}$$

Choosing $(\frac{1}{a} + \frac{1}{a'})(1 + \frac{1}{b'}) < \frac{1}{2}$, $4b > a$, and $A > 2bL$, we obtain the oracle inequality as desired.

Remark. There is an error in the proof of Theorem 1 in Bunea et al. (2012): Although $\boldsymbol{X}(\hat{\boldsymbol{B}}-\boldsymbol{B}) = \boldsymbol{X}_{\mathcal{J}\cap\hat{\mathcal{J}}}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}\cap\hat{\mathcal{J}}}+\boldsymbol{X}_{\mathcal{J}^c\cap\hat{\mathcal{J}}}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}^c\cap\hat{\mathcal{J}}}+\boldsymbol{X}_{\mathcal{J}\cap\hat{\mathcal{J}}^c}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}\cap\hat{\mathcal{J}}^c}$, the inequality $\|\boldsymbol{X}_{\mathcal{J}\cap\hat{\mathcal{J}}}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}\cap\hat{\mathcal{J}}}\|_F^2 + \|\boldsymbol{X}_{\mathcal{J}^c\cap\hat{\mathcal{J}}}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}^c\cap\hat{\mathcal{J}}}\|_F^2 + \|\boldsymbol{X}_{\mathcal{J}\cap\hat{\mathcal{J}}^c}(\hat{\boldsymbol{B}}-\boldsymbol{B})_{\mathcal{J}\cap\hat{\mathcal{J}}^c}\|_F^2 \lesssim \|\boldsymbol{X}(\hat{\boldsymbol{B}}-\boldsymbol{B})\|_F^2$ does *not* hold in general. Our approach of proving Theorem 3 is different and avoids the issue.

6.4 Proof of Theorem 4

Let $h(\boldsymbol{B}) = 1/(mn - AP_0(\boldsymbol{B})/\sigma^2)$. From the definition of $\hat{\boldsymbol{B}}$, we have

$$
\begin{aligned}
&\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 \\
\leq &\|\boldsymbol{E}\|_F^2 \left(\frac{h(\boldsymbol{B}^*)}{h(\hat{\boldsymbol{B}})} - 1\right) + 2\langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\rangle \\
= &\frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2 - AP_0(\boldsymbol{B}^*)} P_0(\boldsymbol{B}^*) - \frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2 - AP_0(\boldsymbol{B}^*)} P_0(\hat{\boldsymbol{B}}) + 2\langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\rangle \\
\leq &\frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2 - AP_0(\boldsymbol{B}^*)} P_0(\boldsymbol{B}^*) - \frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2} P_0(\hat{\boldsymbol{B}}) + 2\langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\rangle
\end{aligned}
$$

Section 6.3 gives the detailed treatment of $\langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\rangle$. Here, we use a high-probability result. For example, for term I in (28), Lemma 5 shows that the following inequality holds for any $a, b, a' > 0$

$$
\langle \boldsymbol{E}, \mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}\rangle \leq (\frac{1}{a} + \frac{1}{a'})\|\mathbf{P}_{\mathcal{J}}\boldsymbol{X}\boldsymbol{\Delta}\mathbf{P}_{rs}\|_F^2 + bLP_o(J, r),
$$

with probability at least $1 - \sum_{J=1}^{p}\sum_{r=1}^{m\wedge J} C\exp\{-c((2\sqrt{b/a}-1)^2 LP_o(J,r)/\sigma^2)\}$, where L is a constant. The probability of the complement is further bounded by $\sum_{J=1}^{p} C\exp\{-c(q \wedge J + J\log(ep/J))\} \leq C\exp(-c\log p)$, noticing that $J\log(ep/J) \geq \log p + \epsilon J$ for some $\epsilon > 0$ and any $J \geq 1$. Hence, following the lines of Section 6.3, we obtain an upper bound of the stochastic term for any given $a, b, a' > 0$

$$
2\langle \boldsymbol{E}, \boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\rangle \leq 2(\frac{1}{a} + \frac{1}{a'})\|\boldsymbol{X}\hat{\boldsymbol{B}} - \boldsymbol{X}\boldsymbol{B}^*\|_F^2 + 4bL\{P_0(\hat{\boldsymbol{B}}) + P_0(\boldsymbol{B}^*)\},
$$

with probability at least $1 - Cp^{-c}$ for some $c, C > 0$.

Let γ, γ' be constants satisfying $0 < \gamma < 1, \gamma' > 0$. On the event $\mathcal{E} = \{(1-\gamma)mn\sigma^2 \le \|\boldsymbol{E}\|_F^2 \le (1+\gamma')mn\sigma^2\}$, we have

$$\begin{aligned}&\frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2 - AP_0(\boldsymbol{B}^*)}P_0(\boldsymbol{B}^*) - \frac{A\|\boldsymbol{E}\|_F^2}{mn\sigma^2}P_0(\hat{\boldsymbol{B}})\\ \le&\frac{(1+\gamma')AA_0}{A_0 - A}P_0(\boldsymbol{B}^*) - (1-\gamma)AP_0(\hat{\boldsymbol{B}}).\end{aligned}$$

From Laurent and Massart (2000b), $\mathcal{E}^c$ occurs with probability at most $C'\exp(-c'mn)$ (with c', C' dependent on γ, γ').

When A_0 is large enough (e.g. $A_0 > 4L$), choosing a, a', b, A such that $(\frac{1}{a} + \frac{1}{a'}) < \frac{1}{2}$, $4b > a$, and $4bL \le (1-\gamma)A$ completes the proof.

6.5 Proof of Theorem 5

Lemma 7. *Given any $\boldsymbol{Y} \in \mathbb{R}^{p\times m}$, $\hat{\boldsymbol{B}} = \vec{\Theta}^{\#}(\boldsymbol{Y}; \mathfrak{q}, \eta)$ is a globally optimal solution to*

$$\min_{\boldsymbol{B}} \frac{1}{2}\|\boldsymbol{Y} - \boldsymbol{B}\|_2^2 + \frac{\eta}{2}\|\boldsymbol{B}\|_2^2 =: f(\boldsymbol{B}; \eta) \quad \text{s.t. } J(\boldsymbol{B}) \le \mathfrak{q}. \tag{33}$$

Let $\mathcal{J} \subset [p]$ with $|\mathcal{J}| = \mathfrak{q}$. Assuming $\boldsymbol{B}[\mathcal{J}^c,] = \boldsymbol{0}$, we get the optimal solution $\hat{\boldsymbol{B}}$ with $\hat{\boldsymbol{B}}[\mathcal{J},] = \boldsymbol{Y}[\mathcal{J},]/(1+\eta)$. It follows that $f(\hat{\boldsymbol{B}}; \eta) = \frac{1}{2}\|\boldsymbol{Y}\|_F^2 + \frac{1}{2(1+\eta)}\|\boldsymbol{Y}[\mathcal{J}^c,]\|_F^2$. Hence the group quantile thresholding $\vec{\Theta}^{\#}(\boldsymbol{Y}; \mathfrak{q}, \eta)$ yields a global minimizer.

It remains to study the minimization of $g(\boldsymbol{S}) = \|\boldsymbol{Y}' - \boldsymbol{X}\boldsymbol{S}\|_F^2/(2K) + \eta\|\boldsymbol{S}\|_F^2/2$ s.t. $\|\boldsymbol{S}\|_{2,0} \le \mathfrak{q}$, where $\boldsymbol{Y}' = \boldsymbol{Y}\boldsymbol{V}$. Construct a surrogate function $G(\boldsymbol{S}, \tilde{\boldsymbol{S}}) = g(\tilde{\boldsymbol{S}}) + \langle(\boldsymbol{I} - \boldsymbol{X}^T\boldsymbol{X}/K)(\boldsymbol{S} - \tilde{\boldsymbol{S}}), \boldsymbol{S} - \tilde{\boldsymbol{S}}\rangle$. The rest of the proof follows the same lines as She (2009). The details are omitted.

References

Aldrin, M. (1996). Moderate projection pursuit regression for multivariate response data. *Comput. Stat. Data Anal.*, 21:501–531.

Anderson, T. W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *Annals of Mathematical Statistics*, 22:327–351.

Bunea, F., She, Y., and Wegkamp, M. (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *Annals of Statistics*, 39:1282–1309.

Bunea, F., She, Y., and Wegkamp, M. (2012). Joint variable and ank selection for parsimonious estimation of high dimensional matrices. *Annals of Statistics*. to appear.

Cai, T. T., Ma, Z., and Wu, Y. (2013). Sparse pca: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110.

Candès, E. J. and Plan, Y. (2011). Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359.

Chen, J. and Chen, Z. (2008). Extended Bayesian information criterion for model selection with large model space. *Biometrika*, 95:759–771.

Chen, L. and Huang, J. Z. (2012). Sparse reduced-rank regression for simultaneous dimension reduction and variable selection. *Journal of the American Statistical Association*, 107:1533–1545.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96:1348–1360.

Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society Series B*, 70(5):849–911.

Foster, D. P. and George, E. I. (1994). The risk inflation criterion for multiple regression. *Annals of Statistics*, 22(4):19471975.

Georghiades, A., Belhumeur, P., and Kriegman, D. (2001). From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660.

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24:417–441, 498–520.

Izenman, A. (1975). Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis*, 5:248–262.

Izenman, A. (2008). *Modern Multivariate. Statistical Techniques: Regression, Classification and Manifold Learning*. Springer, New York.

Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693.

Jolliffe, I. T., Trendafilov, N. T., and Uddin, M. (2003). A modified principal component technique based on the lasso. *Journal of Computational and Graphical Statistics*, 12:531–547.

Journée, M., Nesterov, Y., Richtárik, P., and Sepulchre, R. (2010). Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.*, 11:517–553.

Laurent, B. and Massart, P. (2000a). Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.*, 28(5):1302–1338.

Laurent, B. and Massart, P. (2000b). Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338.

Lee, K., Ho, J., and Kriegman, D. (2005). Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 27(5):684–698.

Lounici, K., Pontil, M., Tsybakov, A. B., and van de Geer, S. (2011). Oracle inequalities and optimal inference under group sparsity. *Annals of Statistics*, 39:2164–2204.

Luenberger, D. and Ye, Y. (2008). *Linear and Nonlinear Programming*. Springer, third edition.

Ma, Z. (2013). Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801.

Mackey, L. (2008). Deflation methods for sparse pca. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1014–1024, Vancouver, Canada.

Owen, A. B. (2007). A robust hybrid of lasso and ridge regression. *Prediction and Discovery (Contemporary Mathematics)*, 443:59–71.

Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642.

Recht, B., Fazel, M., and Parrilo, P. A. (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501.

Reinsel, G. and Velu, R. (1998). *Multivariate Reduced-Rank Regression: Theory and Applications*. Springer, New York.

She, Y. (2009). Thresholding-based iterative selection procedures for model selection and shrinkage. *Electronic Journal of Statistics*, 3:384–415.

She, Y. (2012). An iterative algorithm for fitting nonconvex penalized generalized linear models with grouped predictors. *Computational Statistics and Data Analysis*, 9:2976–2990.

She, Y. (2013). Reduced rank vector generalized linear models for feature extraction. *Statistics and Its Interface*, 6:197–209.

Shen, H. and Huang, J. (2008). Sparse principal component analysis via regularized low rank matrix approximation. *Journal of Multivariate Analysis*, 99:1015–1034.

Szarek, S. (1982). Nets of grassmann manifold and orthogonal groups. In *Proceedings of Banach Spaces Workshop*, pages 169–185. University of Iowa Press.

Talagrand, M. (2005). *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes*. Springer Monographs in Mathematics. Springer.

Witten, D., Tibshirani, R., and Hastie, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Journal of Multivariate Analysis*, 10(3):515–534.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *JRSSB*, 68:49–67.

Zhang, C.-H. (2010a). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942.

Zhang, C.-H. and Zhang, T. (2012). A general theory of concave regularization for high dimensional sparse estimation problems. *Statist. Sci.*, 27(4):576–593.

Zhang, T. (2010b). Analysis of multi-stage convex relaxation for sparse regularization. *J. Mach. Learn. Res.*, 11:1081–1107.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *JRSSB*, 67(2):301–320.

Zou, H., Hastie, T., and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15:265–286.